Approximate Safety Verification and Control of Partially Observable Stochastic Hybrid Systems

Kendra Lesser, *Member, IEEE*, Meeko Oishi, *Member, IEEE*

arXiv:1410.8054v3 [cs.SY] 6 Jul 2015

Abstract

Assuring safety in discrete time stochastic hybrid systems is particularly difficult when only noisy or incomplete observations of the state are available. We first review a formulation of the probabilistic safety problem under noisy hybrid observations as a dynamic program over an equivalent information state. Two methods for approximately solving the dynamic program are presented. The first method approximates the hybrid system as an equivalent finite state Markov decision process, so that the information state is a probability mass function. The second approach approximates an indicator function over the safe region using radial basis functions, to represent the information state as a Gaussian mixture. In both cases, we discretize the hybrid observation process and generate a sampled set of information states, then use point-based value iteration to under-approximate the safety probability and synthesize a suboptimal control policy. We obtain error bounds and convergence results in both cases, assuming switched affine dynamics and additive Gaussian noise on the continuous states and observations. We compare the performance of the finite state and Gaussian mixture approaches on a simple numerical example.

I. Introduction

Safety critical systems, such as aircraft, satellites, and electricity grids, often rely on sensors to measure their state and their environment. The true state of the system may not be measurable, or may be corrupted by noise, as quantified by a so-called observation process. The controller must choose actions based only on the information contained in the observation process. The nature of safety critical systems dictates a need for formal methods to accurately assess the system's ability to meet rigorous safety requirements, and also to synthesize controllers that guarantee performance to a desired level (correct by design). It is therefore paramount that the controller exploit information from the observation process, to obtain theoretical safety guarantees that are as accurate as possible.

Reachability analysis, which determines whether a system's state remains within a given safe region and/or reaches a target set within some time horizon, has been used extensively as a tool for verification and controller synthesis for hybrid systems [1], [2], [3] and extended to stochastic hybrid systems (SHS) [4], [5]. There has been little focus,

K. Lesser is with the Department of Computer Science, University of Oxford, OX1 3QD, United Kingdom (e-mail: kendra.lesser@cs.ox.ac.uk).

M. Oishi is with the Department of Electrical and Computer Engineering, University of New Mexico, Albuquerque, NM, USA (e-mail: oishi@unm.edu).

This material is based upon work supported by the National Science Foundation under Grant Number CMMI-1254990, CNS-1329878, and CMMI-1335038 and by the Air Force Office of Scientific Research through a Summer Faculty Fellowship. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

however, on reachability analysis for partially observable SHS. While there has been some work on deterministic hybrid systems with hidden modes [6] or uncertain systems with imperfect information on a partial order [7], reachability analysis of a partially observable SHS has been approached only theoretically [8], [9]. This, along with our previous work [10] provides the first computational results for both controller synthesis and verification of safety specifications for partially observable SHS.

Existing computational results for reachability analysis of fully observable SHS are also limited. The safety problem for a discrete time SHS (DTSHS), which considers only whether the state of the system can be controlled to remain within a safe region of the state space, can be formulated as a multiplicative cost stochastic optimal control problem [4], and solved in the same manner as a Markov decision process (MDP). Unfortunately, solutions via dynamic programming [11] require evaluation of a value function over all possible states, which is infinite when those states are continuous. Discretization procedures can be employed to impose a finite number of states, as in [12] and [13], which present rigorous uniform and adaptive gridding methods for verification of DTSHS. Similarly, approximate abstractions of the original stochastic model to an equivalent system that has the same properties are presented in [14], [15], and [16]. Even so, current applications are limited to those with only a few discrete and continuous states.

The safety problem for a partially observable DTSHS (PODTSHS) can similarly be formulated as a partially observable MDP (POMDP). However, POMDPs are plagued by dimensionality on an even greater scale than MDPs. The common approach to solving POMDPs is to replace the growing history of observations and actions by a sufficient statistic, often called the belief state, which, for a POMDP with an additive cost function, is the distribution of the current state conditioned on all past observations and actions [11]. This belief state is treated as the perfectly observed true state, and MDP solution methods can then be applied. However, with a continuous state space, the belief state is a function defined over an infinite domain, and it is impossible to enumerate over all such functions. Therefore the study of efficient, approximate solutions to POMDPs is essential.

Although finding the solution to a general POMDP is hard [17], many algorithms for approximating solutions to finite state POMDPs have been developed. These mainly rely on point-based value iteration (PBVI) schemes that only consider a subset of the belief space to update the value function used in the dynamic program (for a survey of PBVI algorithms, see [18]). Because the value function is piecewise-linear and convex [19] (and so equivalently represented by a finite set of vectors), sampling from the belief state provides a systematic way of storing a finite subset of those vectors. Such methods must be tailored to continuous state POMDPs because of the dimensionality of the belief state.

Other than discretizing the state space and solving an equivalent finite state POMDP, many existing methods for continuous state POMDPs assume the belief state is Gaussian (e.g. [20], [21]), and represent the belief state in a parameterized form which is then discretized and solved as a discrete state MDP. When the belief state cannot adequately be represented using a single Gaussian, a Gaussian mixture may be used instead. An equivalent point-based algorithm for continuous-state POMDPs using Gaussian mixtures is presented in [22], and demonstrated on a stochastic hybrid system with hidden modes in [23].

The safety problem for a PODTSHS is further complicated because the belief state is not the conditional

distribution of the current state of the system [8], [9], but must also include the distribution of a binary variable that indicates whether the state of the system has remained within a safe region up to the previous time step. This, coupled with the stochastic hybrid system dynamics, makes accurately representing the belief state as a single Gaussian impossible.

We formulate the safety problem for a PODTSHS as a POMDP, and investigate representations of the belief state in either vector or Gaussian mixture form through finite- and continuous-state approximations to the PODTSHS. These representations allow us to exploit point-based methods developed for POMDPs.

This paper extends our previous work [10] in several ways. First, we validate the use of POMDP solution techniques for reachability analysis of a PODTSHS, by demonstrating that the value function is convex and admits a function representation related to the piecewise-linear vector representation of a finite state POMDP. Second, we present a finite state approximation to the DTSHS (presented in [10] without proofs) that allows the belief state to take vector form under certain conditions, and show convergence for the approximation. Third, we preserve the continuity in the hybrid state space through a Gaussian mixture representation for the belief state, and approximate the indicator function that represents the safe region using Gaussian radial basis functions. In this case, we provide an error bound as a function of the integrated error (1-norm in the function space L^1) of the indicator function approximation. Our solution method converges to the true solution from below, using either the finite or continuity-preserving belief state. We demonstrate both approaches on a temperature regulation problem.

The rest of the paper is organized as follows. Section II relates the safety problem for a PODTSHS to optimal control of a POMDP. Section III justifies the use of POMDP solution techniques, and presents the finite and Gaussian mixture approximations to the safety problem for a PODTSHS (as well as error bounds). Section IV describes the use of point-based approximation techniques, through sampling of belief states and discretization of the observations. We present a numerical example in Section V, and concluding remarks and directions for future work in Section VI. All proofs can be found in the Appendix.

II. Background

A. Notation and Probability Overview

A probability space $(\Omega, \mathcal{F}, \mathbb{P})$ consists of a sample space Ω, a σ-algebra $\mathcal{F}$ defined over Ω, and a probability measure $\mathbb{P}$ that assigns probabilities to events in $\mathcal{F}$. For $\Omega = \mathbb{R}^n$, we presume $\mathcal{F} = \mathcal{B}(\mathbb{R}^n)$, the Borel σ-algebra on $\mathbb{R}^n$. The probability measure $\mathbb{P}$ maps elements $B \in \mathcal{B}(\mathbb{R}^n)$ to the interval $[0, 1]$. The density p associated to $\mathbb{P}$ is defined according to the Lebesgue measure as $\mathbb{P}(B) = \int_B p(x)\, dx$.

We denote expected value by $\mathbb{E}$. A probability measure or expected value induced by a control policy π (to be defined later), is $\mathbb{P}^\pi$ or $\mathbb{E}^\pi$, respectively. For a space $\mathcal{X}$, $\mathcal{X}^n = \mathcal{X} \times \ldots \times \mathcal{X}$ is the n-times product space of $\mathcal{X}$. The state of a system at time n, x_n, takes values in state space $\mathcal{X}$. We use $|\cdot|$ to denote absolute value for $x \in \mathbb{R}$ and cardinality for a set Ω. A vector or function norm is denoted with $\|\cdot\|$. Unless otherwise specified, $\|\cdot\| = \|\cdot\|_2$. A random variable $x \sim \mathcal{N}(\mu, \mathcal{P})$ follows a Gaussian distribution with mean μ and covariance $\mathcal{P}$, and $\phi(x; \mu, \mathcal{P})$ represents a Gaussian density with mean μ and covariance $\mathcal{P}$ evaluated at x. Finally, $\langle \cdot, \cdot \rangle$ denotes an inner

product, with $\langle u, v \rangle = u^T v$ for vectors $u, v \in \mathbb{R}^n$, and $\langle f, g \rangle = \int f(x)g(x)\,dx$ or $\langle f, g \rangle = \sum_q \int f(x,q)g(x,q)\,dx$ for appropriately defined functions f, g.

B. Computing Optimal Control Policies for POMDPs

The main results of this paper rely on framing a PODTSHS as a POMDP with a hybrid state space, and drawing upon existing results for finite state POMDPs with additive cost objectives. We therefore present an overview of POMDPs and efficient approximation techniques for their optimal control, and then express a PODTSHS with a safety objective as a POMDP with a multiplicative cost function.

Definition 1. *(POMDP $\mathcal{J}$) A POMDP is a tuple $\mathcal{J} = (\mathcal{S}, \mathcal{U}, \mathcal{Y}, T, Y, R)$ where*

1) *$\mathcal{S}$ is a finite set of states*
2) *$\mathcal{U}$ is a finite set of possible control inputs*
3) *$\mathcal{Y}$ is a finite set of observations*
4) *$T : \mathcal{S} \times \mathcal{S} \times \mathcal{U} \to [0,1]$ is a state transition function that assigns a probability measure to state s_{n+1} given state s_n and control u_n for all n: $T(s_{n+1}|s_n, u_n)$*
5) *$Y : \mathcal{Y} \times \mathcal{S} \times \mathcal{U} \to [0,1]$ is an observation function that assigns a probability measure to observation y_n given state s_n and control u_{n-1} for all n: $Y(y_n|s_n, u_{n-1})$*
6) *$R : \mathcal{S} \to [0,1]$ is an initial probability measure over the state space $\mathcal{S}$: $R(s)$*

The state evolves stochastically and is Markovian (e.g., the state at the next time step depends only on the current state and action). The information available to the controller at time n is $i_n = (u_0, \ldots, u_{n-1}, y_1, \ldots, y_n) \in \mathcal{I}_n = \mathcal{U}^n \times \mathcal{Y}^n$; that is, the controller cannot directly observe the state. The control input at each time step is selected according to a control policy π, which maps the available information at each time n onto $\mathcal{U}$.

Definition 2. *For a POMDP $\mathcal{J}$, a policy π for some time horizon N is a sequence of functions, $\pi = (\pi_0, \ldots, \pi_{N-1})$, such that $\pi_n : \mathcal{I}_n \to \mathcal{U}$.*

We consider non-randomized policies, i.e. ones that assign a single control input to each possible i_n, which are sufficient for the problem we consider [24]. A control policy π induces a probability space $(\Omega, \sigma(\Omega), \mathbb{P}^\pi)$ over the POMDP with state space $\Omega = \mathcal{S}^n \times \mathcal{Y}^n$, σ-algebra $\sigma(\Omega)$ on Ω, and probability measure $\mathbb{P}^\pi$ based on T, Y, R, and π.

The execution of a POMDP is as follows. At time $n = 0$, state s_0 is produced from initial distribution $R : s \sim R(\cdot)$. At each subsequent time $n > 0$, an observation y_n is produced according to $y_n \sim Y(\cdot|s_n, u_{n-1})$, and added to the list of past observations and control inputs to produce $i_n = (u_0, \ldots, u_{n-1}, y_1, \ldots, y_n)$. The control input is chosen according to $u_n = \pi_n(i_n)$, and cost $C(s_n, u_n)$ is accrued. The next state s_{n+1} is then generated according to $s_{n+1} \sim T(\cdot|s_n, u_n)$.

The goal is to minimize the expected sum of costs accrued according to function C over a time horizon N by

optimally choosing control actions according to the policy π.

$$\max_{\pi} \mathbb{E}^{\pi} \left[\sum_{n=0}^{N} C(s_n, \pi_n(i_n)) \right] \tag{1}$$

Equation (1) can be solved using dynamic programming, much like for a Markov decision process [11]. The value function, $V_n(i_n)$, represents the expected sum of costs accrued from time n to N given that i_n has been recorded thus far, and is computed recursively backwards in time. However, since the size of vector i_n increases with n, and is difficult to store, the optimal control input and value function can instead be expressed as a function of a belief state. The belief state is a *sufficient statistic* for the information vector i_n because it condenses all information necessary for making optimal decisions without sacrificing optimality. For an additive cost function C, the belief state is a function that describes the probability of being in state s given all past observations and actions, $b(s_n) = \mathbb{P}[s_n|i_n]$ [11]. By treating the belief state as the true state of the system, (1) can be equivalently solved as an MDP over the belief state. The optimal policy $\pi^* : \mathfrak{B} \to \mathcal{U}$ is hence defined in terms of the belief state, with $\mathfrak{B}$ the space of all beliefs.

The optimal policy can be found using a value function over the belief space,

$$V_n^*(b) = \max_{u \in \mathcal{U}} \left\{ \sum_{s} C(s,u) b(s) + \sum_{y} V_{n+1}^* \left(M_{y,u} b \right) \mathbb{P}(y|b,u) \right\}. \tag{2}$$

The transition operator $M_{y,u}b$ provides the next belief state b_{n+1} given the current observation, action, and belief state,

$$\left(M_{y,u} b \right)(s') = \frac{Y(y|s',u) \sum_{s \in \mathcal{S}} T(s'|s,u) b(s)}{\mathbb{P}(y|b,u)}, \tag{3}$$

and the likelihood of the observation is $\mathbb{P}(y|b,u) = \sum_{s \in \mathcal{S}} b(s) \sum_{s' \in \mathcal{S}} T(s'|s,u) Y(y|s',u)$. Sondik [19] first showed that for a finite horizon $N < \infty$, the value function at each time n is piecewise linear and convex in the belief state, and thus can be expressed as

$$V_n^*(b) = \max_{\alpha_n \in \Gamma_n} \sum_{s} \alpha_n(s) b(s). \tag{4}$$

The functions $\alpha_n \in \mathbb{R}^{|\mathcal{S}|}$, or "$\alpha$-vectors", characterize the value of being in state $s \in \mathcal{S}$ at time n when a specific action u is taken, plus the expected sum of future rewards assuming all subsequent actions are chosen optimally. Because each α-vector is associated with a specific action, selection of the optimal α-vector in (4) defines the optimal policy for belief b at time n. The set Γ_n of α-vectors needed to exactly represent the value function V_n at time n is finite, but grows exponentially, since computing Γ_n requires calculation of $|\mathcal{U}||\Gamma_{n+1}|^{|\mathcal{Y}|}$ α-vectors. An exact α-vector representation is therefore often infeasible, and approximate solutions are required.

We draw upon methods in point-based value iteration (PBVI) [25], [18] because they provide a lower bound to the value function, which is key for our safety verification problem. In the most general PBVI method, a finite subset $B \subset \mathfrak{B}$ is selected, then one α-vector is generated for each belief point $b^i \in B$, $B = (b^0, b^1, \ldots, b^m)$, so that $\tilde{\Gamma}_n = (\alpha_n^0, \alpha_n^1, \ldots, \alpha_n^m)$. The value at some b not necessarily in B can be approximated by $V_n^*(b) \approx \max_{\alpha_n \in \tilde{\Gamma}_n} \sum_s \alpha_n(s) b(s)$ (as compared to (4)) with $\tilde{\Gamma}_n \subset \Gamma_n$, since we presume that for any b in a neighborhood of b^j the same action will likely be optimal [25]. Hence $\tilde{\Gamma}_n$ can be generated recursively from $\tilde{\Gamma}_{n+1}$ without

enumeration over all possible combinations of observations and subsequent α-vectors in $\tilde{\Gamma}_{n+1}$, by using the backup operation

$$backup(b) = \arg\max_{\alpha_n \in \tilde{\Gamma}_n} \sum_{s \in \mathcal{S}} \alpha_n(s) b(s) \tag{5}$$

for each $b \in B$. Essentially, the PBVI algorithm consists of selecting a set of belief points $B \subset \mathfrak{B}$, and repeatedly applying (5) to each element of B. For N finite, the backup operator is applied to B N times.

For PBVI methods that have been extended to POMDPs with a continuous state space and discrete observation and action spaces [22], the α-vectors are replaced by α-functions defined over the continuous space S. Since observations and actions are finite, the number of α-functions remains finite, and the value function is piecewise-linear and convex under the α-function representation.

C. PODTSHS Modeled as a POMDP

A PODTSHS has discrete and continuous states with interacting dynamics that are characterized by stochastic kernels, and observations that may also be hybrid. It can be expressed within the POMDP framework by extending Definition 1.

Definition 3. *(PODTSHS $\mathcal{H}$). A PODTSHS is a class of POMDP, i.e. a tuple $\mathcal{H} = (\mathcal{S}, \mathcal{U}, \mathcal{Y}, T, Y, R)$, in which*

1) *$\mathcal{S} = \mathbb{R}^m \times \mathcal{Q}$ is a hybrid state space with $\mathcal{Q} = \{q_1, q_2, \ldots, q_{N_q}\}$ a finite set of states with cardinality N_q and continuous state dimension m*
2) *$\mathcal{U}$ is a finite set of control inputs affecting both discrete and continuous state transitions*
3) *$\mathcal{Y} = \mathbb{R}^l \times \mathcal{Y}^q$ is a hybrid observation space with $\mathcal{Y}^q \subseteq \mathcal{Q}$ and continuous observation dimension $l \leq m$*
4) *$T : \mathcal{B}(\mathbb{R}^m) \times \mathcal{Q} \times \mathcal{S} \times \mathcal{U} \to [0,1]$ is a stochastic transition kernel that assigns a probability measure over $\mathcal{B}(\mathbb{R}^m) \times \mathcal{Q}$ at time $n+1$ given s_n and u_n for all n: $T(\beta, q | s_n, u_n)$ with $\beta \in \mathcal{B}(\mathbb{R}^m)$*
5) *$Y : \mathcal{B}(\mathbb{R}^l) \times \mathcal{Y}^q \times \mathcal{S} \to [0,1]$ is an observation function that assigns a probability measure over $\mathcal{B}(\mathbb{R}^p) \times \mathcal{Y}^q$ at time n given s_n for all n: $Y(\beta, y^q | s_n)$*
6) *$R : \mathcal{B}(\mathbb{R}^m) \times \mathcal{Q} \to [0,1]$ is an initial probability measure over $\mathcal{B}(\mathbb{R}^m) \times \mathcal{Q}$: $R(\beta, q)$*

Remark 1. *While we presume $\mathcal{U}$ is finite, a continuous or hybrid input set can be approximated as a finite set when computing safety probabilities and the optimal policy.*

The state transition kernel comprises a discrete component T_q that governs mode updates, and a kernel T_x for continuous state transitions. For modeling purposes, we choose to order (6) such that the discrete mode q_n updates first at each time step, and the subsequent mode q_{n+1} influences the evolution of x_n to x_{n+1}.

$$T(\beta, q | s_n, u_n) = T_x(\beta | q_{n+1}, x_n, u_n) T_q(q_{n+1} | s_n, u_n) \tag{6}$$

The functions Y and R are also separated into discrete and continuous components,

$$Y(\beta, y^q | s_n) = Y_x(\beta | s_n) Y_q(y^q | s_n), \tag{7}$$

$$R(\beta, q) = R_x(\beta) R_q(q). \tag{8}$$

Functions T_x, Y_x, and R_x are Borel-measurable stochastic transition kernels over $\mathcal{B}(\mathbb{R}^{(\cdot)})$, and T_q, Y_q, and R_q are standard probability distributions over finite state spaces.

We consider specifically a switched affine system, such that the continuous state x evolves according to

$$x_{n+1} = A(q_{n+1})x_n + g(q_{n+1}, u_n) + v_n. \tag{9}$$

The v_n are independent and identically distributed Gaussian random variables for all n, $v_n \sim \mathcal{N}(0, \mathcal{V})$. The matrix A and function g change according to the mode q_{n+1}. We assume the discrete observations y^q depend only on q, and the observations y^x depend linearly on x, corrupted by additive Gaussian noise w_n, with w_n independent and identically distributed for all n, $w_n \sim \mathcal{N}(0, \mathcal{W})$.

$$y_n^x = C(q_n)x_n + w_n \tag{10}$$

Kernels T_x, and Y_x admit Gaussian densities $\tau_x(x|q_{n+1}, x_n, u_n) = \phi(x; A(q_{n+1})x_n + g(q_{n+1}, u_n), \mathcal{V})$, and $\gamma_x(y^x|s_n) = \phi(y^x; C(q_n)x_n, \mathcal{W})$. We also assume R_x admits a Gaussian density $\rho_x(x) = \phi(x; \mu_0, \mathcal{P}_0)$. For ease of notation, we let $\rho(s) = \rho_x(x)R_q(q)$, $\gamma(y|s) = \gamma_x(y^x|s)Y_q(y^q|q)$, and $\tau(s'|s,u) = \tau_x(x'|q', x, u)T_q(q'|s, u)$ for $s = (x, q)$, $y = (y^x, y^q)$, and $s' = (x', q')$.

We require that the following Lipschitz properties hold, which are guaranteed for γ_x, τ_x, and ρ_x, given that they are Gaussian densities and $\mathcal{U}$ is finite.

$$\begin{aligned}
|\tau_x(x'|q', x, u) - \tau_x(\overline{x}'|q', x, u)| &\leq h_x^{(1)}\|x' - \overline{x}'\| \\
|\tau_x(x'|q', x, u) - \tau_x(x'|q', \overline{x}, u)| &\leq h_x^{(2)}\|x - \overline{x}\| \\
|\gamma_x(y|x, q) - \gamma_x(\overline{y}|x, q)| &\leq h_y^{(1)}\|y - \overline{y}\| \\
|\gamma_x(y|x, q) - \gamma_x(y|\overline{x}, q)| &\leq h_y^{(2)}\|x - \overline{x}\| \\
|T_q(q'|q, x, u) - T_q(q'|q, \overline{x}, u)| &\leq h_q\|x - \overline{x}\|
\end{aligned} \tag{11}$$

We define the maximum values $\phi_v^* = (2\pi)^{-\frac{m}{2}}|\mathcal{V}|^{-\frac{1}{2}}$ and $\phi_w^* = (2\pi)^{-\frac{l}{2}}|\mathcal{W}|^{-\frac{1}{2}}$.

Remark 2. *While we impose assumptions of linearity and additive Gaussian noise in* (9) *and* (10) *to facilitate subsequent derivations, these assumptions can be relaxed in certain cases, which will be highlighted where appropriate.*

D. Safety Problem

We use stochastic optimal control to find both a control policy that maximizes the probability of the state remaining within a safe region of the state space, as well as an estimate of that probability [4]. For a compact Borel set $K \subseteq \mathcal{B}(\mathcal{S})$, terminal time N, and predefined policy π, the objective to optimize is $p_{\text{safe}}^N(\pi, \rho; K) = \mathbb{P}^\pi[s_n \in K \,\forall\, n = 0, \ldots, N|\rho]$.

This objective can be expressed more commonly as an expected value of a multiplicative cost function, since for a random variable X and event A, $\mathbb{P}[X \in A] = \mathbb{E}[\mathbf{1}_A(X)]$, with the indicator function $\mathbf{1}_A(X) = 1$ if $X \in A$ and $\mathbf{1}_A(X) = 0$ otherwise, as shown in [4]:

$$p_{\text{safe}}^N(\pi, \rho; K) = \mathbb{E}^\pi\left[\prod_{n=0}^{N} \mathbf{1}_K(s_n)\middle|\rho\right]. \tag{12}$$

The maximal safety probability and optimal policy π^* are given by

$$p^N_{\text{safe}}(\rho; K) = \sup_{\pi \in \Pi} p^N_{\text{safe}}(\pi, \rho; K), \tag{13}$$

$$\pi^* = \arg\sup_{\pi \in \Pi} \left\{ p^N_{\text{safe}}(\pi, \rho; K) \right\}. \tag{14}$$

In the fully observable case, [4] gives a dynamic programming formulation for optimizing (12), which returns both the maximal safety probability and optimal policy. We would like to take a similar approach to find both (13) and (14). Formally, we would like to solve the following problem.

Problem 1. *For a PODTSHS $\mathcal{H} = (\mathcal{S}, \mathcal{U}, \mathcal{Y}, T, Y, R)$ with a safe set $K \in \mathcal{B}(\mathcal{S})$ and time horizon $N < \infty$, we wish to*

1) *Compute the maximal probability* (13) *of remaining within K for N time steps.*
2) *Compute the optimal policy π^* given by* (14)*.*

If the maximal probability and optimal policy cannot be computed exactly (which is quite likely [17]), an approximation that produces a suboptimal policy and lower bound on the maximal safety probability is desired.

III. Reformulation using a POMDP

We exploit the PBVI method to solve Problem 1, by transforming Problem 1 into an optimal control problem for a POMDP. Hence we first show the safety problem for $\mathcal{H}$ can be reduced to a dynamic program, despite a non-standard belief state. We then show that the α-functions and belief states can be approximately represented in closed form and that finite collections of each may be generated and used to approximate (13), similar to a point-based POMDP solver.

We present two approximations of Problem 1 for the PODTSHS $\mathcal{H}$: The first discretizes $\mathcal{S}$ to produce a finite state POMDP, and the second preserves continuity in $\mathcal{S}$ by using a Gaussian mixture approach, thus characterizing the PODTSHS by a collection of weights, means, and covariances.

A. Validity of α-Function Representation

The multiplicative nature of the cost function for the safety problem (12) renders the belief state for an additive cost POMDP inapplicable, and we derived a different sufficient statistic $\eta = (\eta_0, \ldots, \eta_N)$ for Problem 1 in [9]. This sufficient statistic produces a modified conditional distribution of the current state that includes the probability that all past states are in the safe set K.

$$\eta_n(\rho, i_n) = \mathbb{P}^\pi[x_n \in \beta, q_n = q, s_0, \ldots, s_{n-1} \in K \mid \rho, i_n] \tag{15}$$

We define the *information state* as the function $\sigma_n(x_n, q_n) \in \Sigma_n \subseteq L^1$ (where L^1 is the space of integrable functions) associated with the probability distribution produced by $\eta_n(\rho, i_n)$, so that $\int_\beta \sigma_n(z, q)\, dz = \eta_n(\rho, i_n)$ for all $q \in \mathcal{Q}$, $\beta \in \mathcal{B}(\mathbb{R}^m)$. Note that the information state is distinct from the belief state (e.g. the conditional

distribution of the current state). The information state updates recursively with a bounded linear operator Φ (for proof see [9])

$$\begin{aligned} \sigma_0 &= \rho \\ \sigma_n &= \Phi_{y_n, u_{n-1}} \sigma_{n-1} \end{aligned} \tag{16}$$

where $\Phi_{y,u}\sigma$ is given by

$$\left(\Phi_{y,u}\sigma\right)(s') = \frac{1}{p(y|\sigma,u)} \gamma(y|s',u) \int_K \tau(s'|s,u)\sigma(s)\, ds. \tag{17}$$

In comparing (3) to (17), the latter integrates over the compact hybrid set K, as opposed to a summation over finite set $\mathcal{S}$.

We define a dynamic programming recursion over σ as

$$\begin{aligned} V_N^*(\sigma_N) &= \langle \sigma_N, \mathbf{1}_K \rangle \\ V_n^*(\sigma_n) &= \max_{u \in \mathcal{U}} \mathbb{E}^\pi \left[V_{n+1}^*(\Phi_{y,u}\sigma_n) \right] \end{aligned} \tag{18}$$

with solution $V_0^*(\rho) = p_{\text{safe}}^N(\rho; K)$. The optimal policy is $\pi^* = (\pi_0^*, \ldots, \pi_{N-1}^*)$, with $\pi_n^*(\sigma_n) = \arg\max_{u \in \mathcal{U}} V_n^*(\sigma_n)$ for all $n \in [0, N]$.

Lemma 1. *For any n, the value function* (18) *can be written as*

$$V_n^*(\sigma_n) = \sup_{\alpha_n \in \Gamma_n} \langle \alpha_n, \sigma_n \rangle.$$

Lemma 1 leads to the following representation of the value function evaluated at some $\sigma \in \Sigma$.

$$V_n^*(\sigma) = \max_{u \in \mathcal{U}} \int_{\mathcal{Y}} \langle \alpha_{y,u,\sigma}, \sigma_n \rangle \, dy, \tag{19}$$

$$\alpha_{y,u,\sigma}(s) = \int_{\mathcal{S}} \alpha_{n+1}^{*(y)}(s') \gamma(y|s') \tau(s'|s,u) \mathbf{1}_K(s)\, ds' \tag{20}$$

with $*(y)$ used to denote the observation-dependent optimal α-function in Γ_{n+1} that maximizes $\langle \alpha_{n+1}, \Phi_{y,u}\sigma_n \rangle$ (the value function $V_{n+1}(\sigma_{n+1})$) for a specific observation y. The set of α-functions at time n is

$$\Gamma_n = \bigcup_{\sigma \in \Sigma} \left\{ \int_{\mathcal{Y}} \alpha_{y,u^*,\sigma}\, dy \right\} \tag{21}$$

with u^* the control inputs chosen according to (49).

Lemma 2. *The value function* (18) *is convex in σ for all $n = 0, \ldots, N$, σ_n^1, $\sigma_n^2 \in L^1$ and $0 \le \lambda \le 1$:*

$$V_n^*(\lambda\sigma_n^1 + (1-\lambda)\sigma_n^2) \le \lambda V_n^*(\sigma_n^1) + (1-\lambda) V_n^*(\sigma_n^2).$$

Lemmas 1 and 2 show that the value function (18) is convex and admits an α-function representation, hence $\mathcal{H}$ is amenable to POMDP solution techniques. However, we cannot use Lemma 1 to solve Problem 1 directly, since Γ_n is not finite and the α-functions and information states have no common structure.

B. Finite State Approximation

We first consider a finite state POMDP [12], whose solution converges to the true safety probability (13) and optimal policy (14). The state space $\mathcal{S}$ is discretized to obtain a vector representation of α and σ. The observation space is unchanged (i.e. hybrid), because the set of observations only affects the finiteness of sets Γ_n and Σ_n. We defer discussion of producing finite collections of Γ_n and Σ_n to section IV.

Given safe set $K \in \mathcal{B}(\mathcal{S})$, let $K = \bigcup_{q\in\mathcal{Q}} K_q \times \{q\}$. Denote $\lambda = \max_{q\in\mathcal{Q}} \mathcal{L}(K_q)$, the finite Lebesgue measure of $K_q \subset \mathbb{R}^m$. Each K_q is partitioned into a finite number of subsets, so that $K_q = \bigcup_{i=1}^{m_q} K_{i,q}$, with $K_{i,q}$ pairwise disjoint (i.e. $K_{i,q} \cap K_{j,q} = \emptyset$ for all $i \neq j$). Finally, let $\delta^x_{i,q}$ be the diameter of partition $K_{i,q}$ so that $\delta^x_{i,q} = \sup\{\|x - \overline{x}\| : x, \overline{x} \in K_{i,q}\}$, with $\delta^x = \max_{i,q} \delta^x_{i,q}$.

The partition of K is denoted by $\mathcal{G}^s = \bigcup_{i=1,\ldots,m_q,q\in\mathcal{Q}} \mathcal{G}^s_{i,q}$ with $\mathcal{G}^s_{i,q} = K_{i,q} \times \{q\}$. Each element $\mathcal{G}^s_{i,q}$ has a representative point $(x^{i,q}, q)$ and the set $K_\delta = \{(x^{i,q}, q) : i = 1, \ldots, m_q, q \in \mathcal{Q}\}$ is the discrete representation of K. We do not consider here how the points $(x^{i,q}, q)$ are selected, but an example is provided in Section V. The function $\xi : K \to K_\delta$ maps a state $s \in \mathcal{G}^s_{i,q}$ to its representative point $(x^{i,q}, q)$ and the function $\Xi : K_\delta \to K$ is the set-valued map from discrete point $(x^{i,q}, q)$ to its corresponding set $\mathcal{G}^s_{i,q}$. We will abuse notation slightly and interchangeably write $\xi(s) = (x^{i,q}, q)$ and $\xi(x) = x^{i,q}$, since q is mapped to itself, and hence the property of interest is the mapping from x to $x^{i,q}$. The discrete state space is defined as $Z_\delta = K_\delta \bigcup \{\psi_s\}$, with ψ_s a single variable that represents all states $s \in \mathcal{S}\backslash K$.

Definition 4. *(POMDP approximation to PODTSHS, $\hat{\mathcal{H}}$). The POMDP approximation is a tuple $\hat{\mathcal{H}} = (Z_\delta, \mathcal{Y}, \mathcal{U}, \tau_\delta, Y, \rho_\delta)$ where*

1) *Z_δ is a finite set of discrete states*
2) *$\mathcal{Y}$ is as defined in Definition 3*
3) *$\mathcal{U}$ is as defined in Definition 3*
4) *$\tau_\delta : Z_\delta \times \mathcal{U} \times Z_\delta \to [0, 1]$ is a discrete state transition function that assigns probabilities to elements of Z_δ*
5) *$Y : \mathcal{B}(\mathbb{R}^l \times \mathcal{Y}^q \times Z_\delta \to [0, 1]$ is as defined in Definition 3, but is conditioned only on values $z \in Z_\delta$ rather than $s \in \mathcal{S}$*
6) *$\rho_\delta : Z_\delta \to [0, 1]$ is a function that assigns probabilities to elements of Z_δ at time zero*

We define the transition function as

$$\tau_\delta(z'|z, u) = \begin{cases} T(\Xi(z')|z, u), & \text{if } z' \in K_\delta \text{ and } z \in K_\delta \\ 1 - \sum_{\overline{z}\in K_\delta} T(\Xi(\overline{z})|z, u), & \text{if } z' = \psi_s,\ z \in K_\delta \\ 1, & \text{if } z' = \psi_s \text{ and } z = \psi_s \\ 0, & \text{if } z' \in K_\delta \text{ and } z = \psi_s \end{cases} \tag{22}$$

with $\sum_{z'\in Z_\delta} \tau_\delta(z'|z, u) = 1$, and the initial distribution ρ_δ on Z_δ as

$$\rho_\delta(z) = \begin{cases} R(\Xi(z)), & \text{if } z \in K_\delta \\ 1 - \sum_{\overline{z}\in K_\delta} R(\Xi(\overline{z})) & \text{if } z = \psi_s \end{cases} \tag{23}$$

The discrete probability space is $(\Omega_\delta, \sigma(\Omega_\delta), \mathbb{P}_\delta^{\pi_\delta})$ with $\Omega_\delta = Z_\delta^{N+1} \times \mathcal{Y}^N$, $\sigma(\Omega_\delta)$ the σ-algebra on Ω_δ, and $\mathbb{P}_\delta^{\pi_\delta}$ the probability measure uniquely defined by ρ_δ, Y, τ_δ, and a control policy $\pi_\delta = (\pi_0^\delta, \ldots, \pi_{N-1}^\delta)$, $\pi_n^\delta : \Sigma_{n,\delta} \to \mathcal{U}$, with $\Sigma_{n,\delta}$ the set of all information states $\sigma_{n,\delta}$ defined on Z_δ at time n.

We further define the operator $\Phi_{y,u}^\delta$ and the intermediate vector $\alpha_{y,u,\sigma_\delta}^\delta$ for any $y \in \mathcal{Y}$, $u \in \mathcal{U}$, $z', z \in Z_\delta$ as

$$\left(\Phi_{y,u}^\delta \sigma_\delta\right)(z') = \frac{1}{p(y|\sigma_\delta, u)} \gamma(y|z') \sum_{z \in K_\delta} \tau_\delta(z'|z,u)\sigma_\delta(z) \tag{24}$$

$$\alpha_{y,u,\sigma_\delta}^\delta(z) = \sum_{z' \in K_\delta} \alpha_{n+1,\delta}^{*(y)}(z')\gamma(y|z')\tau_\delta(z'|z,u)\mathbf{1}_{K_\delta}(z). \tag{25}$$

The safety problem for $\hat{\mathcal{H}}$ is to find $p_{\text{safe},\delta}^N(\rho_\delta; K_\delta) = \sup_{\pi_\delta \in \Pi_\delta} \mathbb{P}^{\pi_\delta}\left[z_n \in K_\delta, \forall\, n = 0, \ldots, N | \rho_\delta\right]$, which is solved by formulating the information state $\sigma_{n,\delta}$ and the value function $V_{n,\delta}^* : \sigma_{n,\delta} \to [0,1]$ for $n = 0, \ldots, N$.

The discrete information state represents a probability mass function over Z_δ, and can be expressed as an integral over an equivalent density (just as $\tau_\delta(z'|z,u) = T(\Xi(z')|z,u)$).

$$\sigma_{n,\delta}(z) = \begin{cases} \int_{\Xi(z)} \hat{\sigma}_n(s)\, ds, \text{ if } z \in K_\delta \\ \int_{\mathcal{S}\setminus K} \hat{\sigma}_n(s)\, ds, \text{ if } z = \psi_s \end{cases} \tag{26}$$

with $\hat{\sigma}_n(s)$ given by

$$\hat{\sigma}_n(s') = \begin{cases} \rho(s'), & \text{if } n = 0 \\ \left(\hat{\Phi}_{y,u}\hat{\sigma}_{n-1}\right)(s') = \frac{1}{p(y|\hat{\sigma}_{n-1},u)}\gamma(y|\xi(s')) \int_K \tau(s'|\xi(s),u)\hat{\sigma}_{n-1}(s)\, ds, & \text{if } n > 0 \end{cases} \tag{27}$$

This can be verified by substituting the expression for τ_δ in terms of τ into (24) and using an induction argument.

The value function is

$$\begin{aligned} V_{N,\delta}^*(\sigma_{N,\delta}) &= \sum_{z \in K_\delta} \sigma_{N,\delta}(z) \\ V_{n,\delta}^*(\sigma_{n,\delta}) &= \max_{u \in \mathcal{U}} \int_{\mathcal{Y}} V_{n+1,\delta}^*(\Phi_{y,u}^\delta \sigma_{n,\delta}) \mathbb{P}_\delta(dy|\sigma_{n,\delta}, u) \end{aligned} . \tag{28}$$

The maximum probability of remaining within K_δ over N time steps is

$$p_{\text{safe},\delta}^N(\rho_\delta; K_\delta) = V_{0,\delta}^*(\rho_\delta). \tag{29}$$

The safety probability for the finite state approximation $\hat{\mathcal{H}}$ converges to the true solution as grid size parameter δ^x tends to zero. To show this, we first describe the error between the continuous information state σ_n and the vector approximation $\sigma_{n,\delta}$.

Theorem 1. *The density $\hat{\sigma}$ defined in* (27) *satisfies*

$$|\sigma_n(s) - \hat{\sigma}_n(s)| \leq \eta_n^\sigma \delta^x$$

for all $s \in \mathcal{S}$, $\sigma_n \in \Sigma_n$, and η_n^σ given by

$$\eta_n^\sigma = \sum_{i=0}^{n-1} c_{1,i} \left(\prod_{j=i+1}^{n-1} c_{2,j} \right),$$

with $c_{1,i} = \min\{\frac{1}{p(y|\sigma_i,u)}, \frac{1}{p(y|\hat{\sigma}_i,u)}\}[\phi_v^* h_y^{(2)} + \phi_w^* h_x^{(2)} + \phi_w^* \phi_v^* h_q]$, $c_{2,j} = \min\{\frac{1}{p(y|\sigma_j,u)}, \frac{1}{p(y|\hat{\sigma}_j,u)}\}\phi_w^* N_q \lambda$.

To prove convergence of the value function $V^*_{n,\delta}$ to V^*_n, we must first show that integration over the infinite spaces $\mathcal{Y}$ and $\mathcal{S}$ results in a bounded solution.

Consider the following two lemmas regarding integration of γ_x and τ_x over unbounded sets $\mathbb{R}^l$ and $\mathbb{R}^m$, respectively.

Lemma 3. *For any $x, \overline{x} \in K_{i,q}$, for all $i = 1, \dots, m_q$, $q \in \mathcal{Q}$, the following holds:*

$$\int_{\mathbb{R}^l} |\gamma_x(y^x|x,q) - \gamma_x(y^x|\overline{x},q)|\, dy^x \leq \left[\beta^y_{1,i,q} h^{(2)}_y + \beta^y_{2,q}\right] \delta^x_{i,q}$$

with $\beta^y_{1,i,q} = \int_{\{y^x : \|y^x - C(q)x\|^2 \leq \lambda^w_1, y^x \in \mathbb{R}^l, x \in K_{i,q}\}} 1\, dy^x$ and $\beta^y_{2,q} = \phi^*_w \|C(q)\|$. In other words, $\beta^y_{1,i,q}$ is the Lebesgue measure of region $K_{i,q}$ mapped to the observation space $\mathbb{R}^l$ via the matrix $C(q)$, and increased by $\sqrt{\lambda^w_1}$ in all directions, with $\lambda^w_1 = \lambda_{\max}(\mathcal{W})$ the largest eigenvalue of $\mathcal{W}$.

A similar result holds for the integral of τ_x over $\mathbb{R}^m$.

Lemma 4. *For any $x, \overline{x} \in K_{i,q}$, for all $i = 1, \dots, m_q$, $q \in \mathcal{Q}$, and any $u \in \mathcal{U}$, $q' \in \mathcal{Q}$, the following holds:*

$$\int_{\mathbb{R}^m} |\tau_x(x'|q',x,u) - \tau_x(x'|q',\overline{x},u)|\, dx \leq \left[\beta^x_{1,i,q} h^{(2)}_x + \beta^x_{2,q}\right] \delta^x_{i,q},$$

$\beta^x_{1,i,q} = \int_{\{x' : \|x' - A(q)x - g(q,u)\|^2 \leq \lambda^v_1, x' \in \mathbb{R}^m, x \in K_{i,q}\}} 1\, dx$ *and* $\beta^x_{2,q} = \phi^*_v \|A(q)\|$, *with λ^v_1 the largest eigenvalue of $\mathcal{V}$.*

In order to show convergence of (29) to (13), we require some additional definitions. First, similarly to $\hat{\sigma}_n$, we define piecewise constant function $\hat{\alpha}_n$ as $\hat{\alpha}_n(s) = \alpha_{n,\delta}(\xi(s))$, so that

$$\hat{\alpha}_n(s) = \int_{\mathcal{S}} \int_{\mathcal{Y}} \hat{\alpha}^{*(y)}_{n+1}(s') \gamma(y|\xi(s')) \times \tau(s'|\xi(s),u) \mathbf{1}_{K_\delta}(\xi(s))\, dy\, ds'.$$

We also define $\tilde{\alpha}_n(s)$ in the same manner as $\hat{\alpha}_n(s)$, except that it is directly related to $\alpha_n(s)$, i.e. uses the same optimal control input u, and the same combination of α_{n+1}-functions (determined by $*(y)$). In other words, $\tilde{\alpha}_n(s)$ is identical to $\alpha_n(s)$ in terms of the optimal policy tree from time n to N, but the values are calculated using $\gamma(y|\xi(s'),u)$ and $\tau(s'|\xi(s),u)$,

$$\tilde{\alpha}^i_n(s) = \int_{\mathcal{S}} \int_{\mathcal{Y}} \tilde{\alpha}^{i(y)}_{n+1}(s') \gamma(y|\xi(s')) \tau(s'|\xi(s),u^i) \mathbf{1}_{K_\delta}(\xi(s))\, dy, \tag{30}$$

for a specific α-function i associated with α^i_n. The superscript i for u^i and $i(y)$ indicates that the same choice of u and combination of $\alpha^j_{n+1}(s)$ are used for both $\alpha^i_n(s)$ and $\tilde{\alpha}^i_n(s)$. A bound on the difference between $\alpha^i_n(s)$ and $\tilde{\alpha}^i_n(s)$ is given in the following lemma.

Lemma 5. *For any $n \in [0, N]$, and any function $\alpha^i_n(s) \in \Gamma_n$, the associated function $\tilde{\alpha}^i_n(s)$ defined in* (30) *satisfies*

$$|\alpha^i_n(s) - \tilde{\alpha}^i_n(s)| \leq (N-n) N_q \left[\beta^y_1 h^{(2)}_y + \beta^x_1 h^{(2)}_x + \beta^y_2 + \beta^x_2 + h_q\right] \delta^x$$

for all $s \in \mathcal{S}$. The constants β^y_1 and β^x_1 are equal to $\max_{i=1,\dots,m_q, q \in \mathcal{Q}} \beta^y_{1,i,q}$ and $\max_{i=1,\dots,m_q, q \in \mathcal{Q}} \beta^x_{1,i,q}$ from Lemmas 3 and 4, respectively.

We now can show convergence of the approximate safety probability over the discretized state space to the true safety probability.

Theorem 2. *For any time $n \in [0, N]$, and any $\sigma_n \in \Sigma_n$, $\sigma_{n,\delta} \in \Sigma_{n,\delta}$, the error between the value function* (18) *and the value function* (28) *based on the finite state approximation is bounded above by*

$$\left|V_n^*(\sigma_n) - V_{n,\delta}^*(\sigma_{n,\delta})\right| \leq \eta_n^\alpha \delta^x$$

with $\eta_n^\alpha = N_q \lambda \eta_n^\sigma + (N-n) N_q (\beta_1^y h_y^{(2)} + \beta_1^x h_x^{(2)} + \beta_2^y + \beta_2^x)$.

Specifically, the safety probability for PODTSHS $\mathcal{H}$ over time horizon N satisfies

$$\left|p_{\text{safe}}^N(\rho; K) - p_{\text{safe},\delta}^N(\rho_\delta; K_\delta)\right| \leq \left[N_q N(\beta_1^y h_y^{(2)} + \beta_1^x h_x^{(2)} + \beta_2^y + \beta_2^x)\right] \delta^x.$$

Theorem 2 shows that the finite state approximation $\hat{\mathcal{H}}$ provides a means to approximately compute (13) through the safety probability for $\hat{\mathcal{H}}$, (29). As $\delta^x \to 0$, the finite state safety probability (29) converges to the true value (13), and the policy π_δ^* converges to π^*.

Remark 3. *The linearity and additive Gaussian noise assumptions for the continuous state dynamics and observations* (9) *and* (10) *are required only in the proofs of Lemmas 3 and 4, and also in the sense that they guarantee the Lipschitz properties* (11) *used in all proofs. Therefore, nonlinear dynamics and non-Gaussian noise are permissible, so long as* (11) *is satisfied, and an equivalent bound on the integrals in Lemmas 3 and 4 can be derived.*

C. Gaussian Mixture Approximation

We now consider a different approximation by representing the information state and α-functions as Gaussian mixtures, such that each are characterized by a finite set of weights, means, and covariances, dependent on the mode q.

Difficulty arises from the incorporation of the indicator function $\mathbf{1}_K$ in (20) and (17). Integration over the compact set K rather than all of $\mathcal{S}$ violates the preservation of the Gaussian form of σ under operator $\Phi_{y,u}$, and similarly for the α-functions. To preserve the Gaussian mixture structure, we therefore propose a radial basis function (RBF) approximation [27] to the indicator function, using Gaussians as the basis function. For each K_q, we set

$$\mathbf{1}_{K_q}(x) \approx \sum_{i=1}^{I_q} w_i(q) \phi(x; \mu_i(q), \mathcal{P}_i(q)). \tag{31}$$

Typically, both the centers $\mu_i(q)$ and the number of components I_q are fixed beforehand. The weights and covariances can then be chosen to optimize the approximation. For simplicity we will denote $\phi(x; \mu_i(q), \mathcal{P}_i(q))$ by $\phi_i(x)$. This approximation is valid since the RBFs are dense in L^p for $1 \leq p \leq \infty$ [27], i.e. given any function f in L^p, a weighted combination of RBFs can approximate f to arbitrary accuracy given enough components, and $\mathbf{1}_K$ is in L^1.

However, the discontinuity in $\mathbf{1}_{K_q}$ produces the Gibbs phenomenon at the boundary of K_q in the RBF approximation. Although these oscillations will always exist for a finite number of components, they could possibly be mitigated [28]. The oscillations can be constrained to a smaller region of K (shorter wavelength) with the addition of more components, indicating that the integrated error may be reduced even if the pointwise error is not. Because we are interested only in integrating over K, this works to our advantage.

We define a new operator Φ^g and a new α-function $\alpha^g_{y,u,\sigma}$ by inserting the RBF approximation (31) into (17) and (20), respectively. We do not include the scaling factor $\frac{1}{p(y|\sigma,u)}$ in the operator Φ^g, which is difficult to incorporate into a Gaussian mixture. If the information state only represented the probability distribution over the current state of the system, we could simply ensure the weights all sum to one, but it includes the probability that past states were in K, and hence is not guaranteed to integrate to one. Once we implement PBVI in Section IV, we are only concerned with finding the α-function that maximizes the inner product $\langle \alpha_{n,g}, \sigma_{n,g} \rangle$, which is unaffected by a constant scaling factor, and hence we lose nothing by disregarding it. We also presume continuous observations, as in Section III-B.

$$\left(\Phi^g_{y,u}\sigma_g\right)(s') = \gamma(y|s') \sum_{q\in\mathcal{Q}} \int_{\mathbb{R}^m} \left[\sum_{i=1}^{I_q} w_i(q)\phi_i(x)\right] \tau(s'|s,u)\sigma_g(s)\,dx \tag{32}$$

$$\alpha^g_{y,u,\sigma}(s) = \int_{\mathcal{S}} \alpha^{*(y)}_{n+1,g}(s')\gamma(y|s')\tau(s'|s,u)\,ds' \left[\sum_{i=1}^{I_q} w_i(q)\phi_i(x)\right] \tag{33}$$

We provide two lemmas stating that the operator $\Phi^g_{y,u}$ and the α-function update $\alpha^g_{y,u,\sigma}$ preserve the Gaussian mixture representation of $\sigma_{n,g}$ and $\alpha_{n,g}$ for all n. These lemmas rely on the fact that τ_x and γ_x are Gaussian, and that a product of Gaussian distributions is again a Gaussian. We need some additional assumptions, however, to simplify the derivations and make the following lemmas valid.

Assumption 1. *The transition kernel T_q does not depend on x, i.e. the mode update is dependent only on the previous mode: $T_q(q'|q,u)$.*

Assumption 2. *The matrices $A(q)$ and $C(q)$ in* (9) *and* (10) *are invertible for all $q \in \mathcal{Q}$.*

Without the first assumption, we would need to represent $T_q(q'|s,u)$ by a Gaussian mixture, which cannot be exact since we are approximating a function over discrete states by a continuous function [23]. The second assumption allows us to express $\tau_x(x'|q',x,u)$ and $\gamma_x(y|s)$ as Gaussian densities over x. While these assumptions are quite restrictive, they can be relaxed at a cost of introducing additional error, as discussed later in Remark 4.

Lemma 6. *The operator $\Phi^g_{y,u}$ is closed under Gaussian mixtures, i.e. for σ_g a Gaussian mixture with L components, $\Phi^g_{y,u}\sigma_g$ is also a Gaussian mixture with $N_q I_q L$ components for any $u \in \mathcal{U}$, $y \in \mathcal{Y}$.*

Lemma 7. *The expression for $\alpha^g_{y,u,\sigma}$ is closed under Gaussian mixtures, i.e. if $\alpha^{*(y)}_{n+1,g}$ is a Gaussian mixture with M components, $\alpha^g_{y,u,\sigma}$ is also a Gaussian mixture with $N_q I_q M$ components, for any $u \in \mathcal{U}$, $y \in \mathcal{Y}$, $\sigma \in \Sigma$.*

Lemma 6 implies that we can approximate σ through a Gaussian mixture and use the equivalent update operator $\Phi^g_{y,u}$, hence the Gaussian mixture approximation of σ is $\sigma_{n,g}(x,q) = \sum_{l=1}^{L} w^\sigma_{l,n}(q)\phi(x;\mu^\sigma_{l,n}(q),\mathcal{P}^\sigma_{l,n}(q))$. Similarly, the Gaussian mixture of any α-function is written as $\alpha_{n,g}(x,q) = \sum_{m=1}^{M} w^\alpha_{m,n}\phi(x;\mu^\alpha_{m,n}(q),\mathcal{P}^\alpha_{m,n}(q))$ for all n.

$$\begin{aligned} \sigma_{0,g}(x,q) &= \sigma_0(x,q) = R_q(q)\phi(x;\mu_0;\mathcal{P}_0) \\ \sigma_{n,g}(x,q) &= \sum_{l=1}^{L} w^\sigma_{l,n}(q)\phi(x;\mu^\sigma_{l,n}(q),\mathcal{P}^\sigma_{l,n}(q)) \end{aligned}. \tag{34}$$

Similarly, the Gaussian mixture approximation of any α-function is written:

$$\begin{aligned}\alpha_{N,g}(x,q) &= \sum_{i=1}^{I_q} w_i(q)\phi_i(x) \\ \alpha_{n,g}(x,q) &= \sum_{m=1}^{M} w_{m,n}^{\alpha}\phi(x;\mu_{m,n}^{\alpha}(q),\mathcal{P}_{m,n}^{\alpha}(q))\end{aligned} . \tag{35}$$

The weights, means, and covariances are defined recursively. Their exact representations can be found in the Appendix.

Note that although the Gaussian mixture representation of $\alpha_{y,u,\sigma}^{g}$ has a finite number of components given that the representation of $\alpha_{n+1,g}$ is finite, the actual α-function, $\alpha_{n,g}$, is expressed as the integral of $\alpha_{y,u,\sigma}^{g}$ over $\mathcal{Y}$. Therefore, without the assumption that $\mathcal{Y}$ is finite, $\alpha_{n,g}$ must have an infinite number of components (by breaking the integral over $\mathcal{Y}$ into a summation over regions of size $\Delta \subset \mathcal{Y}$ and taking the limit as $\Delta \to 0$). We take some liberty in overlooking this discrepancy, because it does not affect the proofs in this section. We impose a finite set $\mathcal{Y}$ in Section IV, which makes the Gaussian mixture representation of the α-functions indeed valid, and discuss additional error implications.

The subscript g denotes that we are computing an estimate based on the Gaussian mixture approximation (31). The value function $V_{n,g}^{*}(\sigma_{n,g})$ is described through the recursion

$$\begin{aligned}V_{N,g}^{*}(\sigma_{N,g}) &= \sum_{\mathcal{Q}}\int_{\mathbb{R}^m}\sum_{i=1}^{I_q} w_i(q)\phi_i(x)\sigma_{N,g}(x,q)\,dx \\ V_{n,g}^{*}(\sigma_{n,g}) &= \max_{u\in\mathcal{U}}\int_{\mathcal{Y}} V_{n+1,g}^{*}(\Phi_{y,u}^{g}\sigma_{n,g})\mathbb{P}(dy|\sigma_{n,g},u)\end{aligned} . \tag{36}$$

The safety problem for the Gaussian mixture approximation is defined as

$$p_{\mathrm{safe},g}^{N}(\rho;K) = V_{0,g}^{*}(\rho). \tag{37}$$

Since τ_x, γ_x, and ρ_x are Gaussian, and based on Assumptions 1 - 2 made above, the Gaussian mixture representations of α and σ are exact, aside from the approximation of $\mathbf{1}_K$ using RBFs. To quantify the error incurred from calculating $V_{0,g}^{*}$ as opposed to V_0^{*} (from integration of (31) over $\mathcal{S}$ rather than over K), we define the error

$$\delta^I = \max_{q\in\mathcal{Q}}\left\|\mathbf{1}_{K_q} - \sum_{i=1}^{I_q} w_i(q)\phi_i(x)\right\|_1 . \tag{38}$$

We first analyze the error between σ_n and $\sigma_{n,g}$, which is stated in terms of the 1-norm rather than the ∞-norm, to be consistent with considering the integrated error in (38).

Theorem 3. *The Gaussian mixture approximation $\sigma_{n,g}$ of σ_n satisfies*

$$\|\sigma_n - \sigma_{n,g}\|_1 \le \gamma_n^{\sigma} N_q \delta^I$$

for any $n \in [0,N]$, $y \in \mathcal{Y}$, and $u \in \mathcal{U}$, with $\gamma_n^{\sigma} = \sum_{j=0}^{n-1}(\phi_w^{*})^{j+1}\phi_{\sigma,j}^{*}$ and $\phi_{\sigma,j}^{*} = \max_{l\in 1,\dots,L;q\in\mathcal{Q}}(2\pi)^{-\frac{m}{2}}|\Sigma_{l,j}^{\sigma}(q)|^{-\frac{1}{2}}$.

To show convergence of (37) to (13), we define the function $\tilde{\alpha}^i_{n,g}(s)$ which utilizes the same policy tree as $\alpha^i_n(s)$ for a specific $\alpha^i_n(s) \in \Gamma_n$. This is equivalent to (30), except that $\tilde{\alpha}^i_{n,g}(s)$ is defined by replacing $\mathbf{1}_K(s)$ with (31).

$$\tilde{\alpha}^i_{n,g}(s) = \int_{\mathcal{S}} \int_{\mathcal{Y}} \tilde{\alpha}^{i(y)}_{n+1,g}(s')\gamma(y|s')\tau(s'|s,u^i)\,dy\,ds' \left[\sum_{i=1}^{I_q} w_i(q)\phi_i(x)\right] \tag{39}$$

with u^i the optimal control input associated with $\alpha^i_n(s)$ and $i(y)$ indicating that $\alpha^{i(y)}_{n+1,g}(s)$ is chosen according to the indices selected by $*(y)$ for $\alpha_n(s)$.

Note that the α-functions are no longer guaranteed to be equal to zero outside of K, and also are not guaranteed to be bounded above by one. So long as (31) is of bounded height, however, the α-functions also remain bounded, and we write $\max_{\alpha_{n,g}\in\Gamma_{n,g}} \|\alpha_{n,g}\|_\infty = \overline{\alpha}_n$. We could also adjust the weights in (31) so that the α-functions cannot exceed one, although this may increase the error δ^I.

The following lemma describes the relation between $\alpha_n(s)$ and $\tilde{\alpha}_n(s)$.

Lemma 8. *For any $n \in [0, N]$, and any $\alpha^i_n(s) \in \Gamma_n$, the associated function $\tilde{\alpha}^i_{n,g}(s)$ defined in* (39) *satisfies*

$$\left\|\alpha^i_n - \tilde{\alpha}^i_{n,g}\right\|_1 \leq \left(\sum_{k=n}^{N} (\lambda\phi^*_v)^{N-k}\overline{\alpha}_{N-k+n+1}\right) N_q\delta^I$$

where $\overline{\alpha}_{N+1} = 1$.

Theorem 4. *For any time $n \in [0, N]$, and any $\sigma_n \in \Sigma_n$, $\sigma_{n,g} \in \Sigma_{n,g}$, the error between the value function* (18) *given σ_n and the value function* (36) *given $\sigma_{n,g}$ using the Gaussian mixture approximation is bounded above by*

$$\left|V^*_n(\sigma_n) - V^*_{n,g}(\sigma_{n,g})\right| \leq \gamma^\alpha_n N_q\delta^I$$

with $\gamma^\alpha_n = \sum_{k=n}^{N} (\lambda\phi^*_v)^{N-k}\overline{\alpha}_{N-k+n+1}\phi^*_{\sigma,n}\delta^I + \overline{\alpha}_n\gamma^\sigma_n$.

Specifically, the safety probability for PODTSHS $\mathcal{H}$ over time horizon N satisfies

$$\left|p^N_{\text{safe}}(\rho; K) - p^N_{\text{safe},g}(\rho; K)\right| \leq \gamma^\sigma_0\delta^I.$$

Theorem 4 shows that the convergence of the Gaussian mixture approximation of both σ and the value function depends on the integrated error between the indicator function over K and the RBF approximation (31), rather than the pointwise error. Although the pointwise error may not converge to zero for a finite number of components in the RBF, the integral of the error can be small, as we will show in Section V.

Remark 4. *The linearity and additive Gaussian noise assumptions on the dynamics* (9) *and* (10)*, as well as Assumptions 1 - 2, are used in the Gaussian mixture approximation to ensure that the only error in the value function and information state approximations comes from the approximation of the indicator function $\mathbf{1}_K(s)$ by a Gaussian mixture. Dropping these assumptions requires that we approximate τ_x, T_q, and γ_x by Gaussian mixtures, which is possible but introduces additional error that we have chosen not to consider.*

IV. Approximate Numerical Solution with Lower Bound

A numerical solution of Problem 1 via either a discrete or Gaussian mixture approximation additionally requires sets Γ_n and Σ_n to be finite, whereas we have sets of infinite size because of the uncountable nature of $\mathcal{Y}$. However,

a lower bound on the safety probabilities (29) and (37) can still be obtained, by characterizing the error that results from using $\tilde{\Gamma}_n \subset \Gamma_n$ and $\tilde{\Sigma}_n \subset \Sigma_n$, finite collections of α-functions and information states, respectively.

We again exploit point-based approximation methods described in Section II-B. We examine the generation of subsets of the information states and α-functions, and prove that each guarantees a lower bound to the safety probability of whichever approximation of Section III we choose. In contrast to most point-based solvers, we do not assume a finite set of observations, and hence discretize the observations merely for the computation of the α-functions. Combining belief space sampling with discretized observations assures a lower bound to the safety probability.

A. Sampling from the information space

We characterize the error from using a sampled subset of Σ_n for performing backup operations (as in (5)). Presume that a finite set of information states $\tilde{\Sigma}_n$ has been generated according to one of the many methods available [18]. We generate a finite set $\tilde{\Gamma}_n$ of α-functions, one for each $\sigma_n \in \tilde{\Sigma}_n$. The convexity of the value functions guarantees that the subset $\tilde{\Gamma}_n$ provides a lower bound on V_n^*. Further, we can show that the error between the approximate value function $\tilde{V}_n^*$ (represented by $\tilde{\Gamma}_n$) and the true value function V_n^* (represented by the complete set of α-functions Γ_n) depends on how densely we sample Σ_n.

We define an intermediate value function $\hat{V}_n^*(\sigma_n) = \sup_{\hat{\alpha}_n \in \tilde{\Gamma}_n} \langle \hat{\alpha}_n, \sigma_n \rangle$ that generates $\tilde{\Gamma}_n$ recursively from Γ_{n+1}, i.e. that introduces one point-based backup from the full set Γ_{n+1}. Then $\hat{\alpha}_n$ is written as a function of α_{n+1}^* rather than $\tilde{\alpha}_{n+1}^*$.

The value function $\tilde{V}_n^*$ is formally defined as $\tilde{V}_n^*(\sigma_n) = \sup_{\tilde{\alpha}_n \in \tilde{\Gamma}_n} \langle \tilde{\alpha}_n, \sigma_n \rangle$

$$\begin{aligned} \tilde{\alpha}_n(s) &= \int_{\mathcal{Y}} \int_{\mathcal{S}} \tilde{\alpha}_{n+1}^{*(y)}(s') \gamma(y|s') \tau(s'|s,u) \mathbf{1}_K(s) \, dy \, ds' \\ \tilde{\alpha}_{n+1}^{*(y)}(s') &= \arg \left\{ \sup_{\tilde{\alpha}_{n+1} \in \tilde{\Gamma}_{n+1}} \int_{\mathcal{S}} \tilde{\alpha}_{n+1}(s') \gamma(y|s') \int_K \tau(s'|s,u) \sigma_n(s) \, ds' \, ds \right\} \end{aligned} \tag{40}$$

so that $\tilde{V}_n^*$ is characterized by the finite set $\tilde{\Gamma}_n$ at each time step. We also define an intermediate value function $\hat{V}_n^* = \sup_{\hat{\alpha}_n \in \tilde{\Gamma}_n} \langle \hat{\alpha}_n, \sigma \rangle$ that generates $\tilde{\Gamma}_n$ recursively from Γ_{n+1}, i.e. that introduces one point-based backup from the full set Γ_{n+1}. Then $\hat{\alpha}_n$ is written as a function of α_{n+1}^* rather than $\tilde{\alpha}_{n+1}^*$.

$$\alpha_{n+1}^{*(y)}(s') = \arg \left\{ \sup_{\alpha_{n+1} \in \Gamma_{n+1}} \int_{\mathcal{S}} \alpha_{n+1}(s') \gamma(y|s') \int_K \tau(s'|s,u) \sigma_n(s) \, ds' \, ds \right\}. \tag{41}$$

We let δ^σ denote the maximum Hausdorff distance over n between points in $\tilde{\Sigma}_n$ and points in Σ_n with respect to the metric $\| \cdot \|_1$.

$$\delta^\sigma = \max_n \left\{ \max\{ \sup_{\tilde{\sigma}_n \in \tilde{\Sigma}_n} \inf_{\sigma_n \in \Sigma_n} \|\tilde{\sigma}_n - \sigma_n\|_1, \sup_{\sigma_n \in \Sigma_n} \inf_{\tilde{\sigma}_n \in \tilde{\Sigma}_n} \|\tilde{\sigma}_n - \sigma_n\|_1 \} \right\} \tag{42}$$

In the following, we do not distinguish between the vector and Gaussian mixture representations of σ and α, because the results apply to both cases.

Lemma 9. *For any $n \in [0, N]$ and $\sigma_n \in \Sigma_n$, the error introduced in one iteration of point-based value iteration is at most δ^σ.*

$$|\hat{V}_n^*(\sigma_n) - V_n^*(\sigma_n)| \leq \delta^\sigma$$

We now use Lemma 9 to derive a bound between the true value function and the point-based approximation at any time n.

Theorem 5. *For a set of information states Σ_n, sampled set $\tilde{\Sigma}_n$, and any time $n \in [0, N]$ and $\sigma_n \in \Sigma_n$, the error from using point-based value iteration versus full value iteration is bounded above by*

$$|\tilde{V}_n^*(\sigma_n) - V_n^*(\sigma_n)| \leq (N - n)\delta^\sigma.$$

Thus the error between the point-based approximation and the actual value function is directly proportional to how densely $\tilde{\Sigma}_n$ is sampled, and converges to zero as $\tilde{\Sigma}_n$ approaches Σ_n. The proofs of Lemma 9 and Theorem 5 are a straightforward extension of those appearing in [25], and are omitted.

B. Calculating the α-functions

Over the uncountably infinite space $\mathcal{Y}$, we cannot calculate $\alpha_{y,u,\sigma}$ for all $y \in \mathcal{Y}$, despite a finite set of u and σ. We therefore compute a subset of $\alpha_{y,u,\sigma}$ for the finite set y^i, to approximate α_n as $\alpha_n(s) \approx \sum_{y^i} \alpha_{y^i,u,\sigma}(s)$. We discretize $\mathcal{Y}$ in a similar fashion to the discretization of $\mathcal{S}$ in Section III-B.

However, since $\mathbb{R}^l$ is not compact, we consider an expanded set $\overline{K} = \bigcup_{y^q \in \mathcal{Q}} \overline{K}_{y^q} \supset K$ defined so that the probability of observing a value y for $s \in K$ that is outside of $\overline{K}$ is approximately zero, i.e. $\gamma(\mathcal{Y} \backslash \overline{K} | s \in K) < \epsilon$, $\epsilon \ll 1$. The sets $\overline{K}_{y^q}$ are divided into disjoint subsets $\overline{K}_{i,y^q}$, $\bigcup_{i=1,\dots l_q} \overline{K}_{i,y^q} = \overline{K}_{y^q}$. We also define $\psi_y = \overline{K}^c = \mathcal{Y} \backslash \overline{K}$, such that $\bigcup_{i=1,\dots,l_q} \overline{K}_{i,y^q} \times \{\psi_y\} = \mathbb{R}^n$.

The partition of $\overline{K}$ is denoted $\mathcal{G}^y = \bigcup_{i,y^q} \mathcal{G}^y_{i,y^q}$ with $\mathcal{G}^y_{i,y^q} = \{\overline{K}_{i,y^q} \times y^q : i = 1, \dots l_q, y^q \in \mathcal{Q}\}$. The diameter of partition $\overline{K}_{i,y^q}$ is $\delta^y_{i,y^q} = \sup\{\|y - \overline{y}\| : y, \overline{y} \in \overline{K}_{i,y^q}\}$, with maximum diameter $\delta^y = \max_{i,y^q} \delta^y_{i,y^q}$. Each partition $\mathcal{G}^y_{i,y^q}$ has a representative element (y^{x,i,y^q}, y^q) and a set $\mathcal{Y}_\delta = \{(y^{x,i,y^q}, y^q) : i = 1, \dots, l_q, y^q \in \mathcal{Q}\}$. The function $\theta : \mathcal{Y} \to \mathcal{Y}_\delta$ maps observation $y \in \mathcal{Y}$ to its representative value (y^{x,i,y^q}, y^q); the set-valued function $\Theta : \mathcal{Y}_\delta \to \overline{K}$ maps the point (y^{x,i,y^q}, y^q) to the corresponding set $\mathcal{G}_{i,y^q}$.

The finite observation space is $W_\delta = \mathcal{Y}_\delta \bigcup \{\psi_y\}$. For the finite state approximation, the transition function $\gamma_\delta : W_\delta \times K_\delta \to [0, 1]$ is defined as

$$\gamma_\delta(w|z) = \begin{cases} Y(\Theta(w)|z), & \text{if } w \in \mathcal{Y}_\delta \\ 1 - \sum_{\overline{w} \in \mathcal{Y}_\delta} Y(\Theta(\overline{w})|z), & \text{if } w = \psi_y \end{cases}. \tag{43}$$

For the Gaussian mixture approximation, we define the transition function γ_g in the same fashion as (43), but with

$$Y(\Theta(w)|z) \approx Y_q(y^q|q) \sum_{j=1}^{M_y} c_j \phi_j^y(y_j^{x,i,y^q}; C(q)x, \mathcal{W}) \tag{44}$$

so that the α-functions will also be Gaussian mixtures at each time step. Note that $w = (y^{x,i,y^q}, y^q)$, y_j^{x,i,y^q} is a set of mesh points inside $\mathcal{G}_{i,y^q}$ associated with w, and c_j is a weight proportional to the mesh spacing (determined, e.g., by the trapezoidal rule for numerical integration).

1) Discretized observations for finite state approximation: We use $\tilde{\Gamma}_{n,\delta}$ and $\tilde{V}^*_{n,\delta}$ to denote the approximation using a finite subset of $\Gamma_{n,\delta}$, with the important distinction that the subset is now generated by a finite collection of observations (as opposed to $\tilde{\Sigma}_n$, i.e. we assume here that the set Σ_n is finite).

The value function is then $\sup_{\tilde{\alpha}_{n,\delta}\in\tilde{\Gamma}_{n,\delta}} \sum_{z\in K_\delta} \tilde{\alpha}_{n,\delta}(z)\sigma_{n,\delta}(z)$, where $\tilde{\alpha}_{n,\delta}$ defined as in Section III-B, only with $\gamma(y|z)$ replaced by $\gamma_\delta(w|z)$.

$$\tilde{V}^*_{n,\delta}(\sigma_{n,\delta}) = \sup_{\tilde{\alpha}_{n,\delta}\in\tilde{\Gamma}_{n,\delta}} \sum_{z\in K_\delta} \tilde{\alpha}_{n,\delta}(z)\sigma_{n,\delta}(z),$$

with

$$\begin{aligned} \tilde{\alpha}_{n,\delta}(z) &= \sum_{w\in W_\delta}\sum_{z'\in K_\delta} \tilde{\alpha}^{*(w)}_{n+1,\delta}(z')\gamma_\delta(w|z')\tau_\delta(z'|z,u) \\ \tilde{\alpha}^{*(w)}_{n+1,\delta}(z') &= \arg\left\{\sup_{\tilde{\alpha}_{n+1,\delta}\in\tilde{\Gamma}_{n+1,\delta}} \sum_{z'\in K_\delta} \tilde{\alpha}_{n+1,\delta}(z')\gamma_\delta(w|z')\tau_\delta(z'|z,u)\sigma_{n,\delta}(z)\right\}. \end{aligned} \tag{45}$$

We again define the intermediate value function $\hat{V}^*_{n,\delta}(\sigma_{n,\delta})$, with $\hat{\alpha}_{n,\delta}$ calculated using $\alpha^{*(w)}_{n+1,\delta} \in \Gamma_{n+1,\delta}$ (as opposed to $\tilde{\Gamma}_{n+1,\delta}$) to capture the error introduce in one backup iteration using discretized observations.

We can then bound the error introduced in one iteration of approximating the α-vectors through discretized observations.

Lemma 10. *For any time $n \in [0, N]$ and $\sigma_{n,\delta} \in \Sigma_{n,\delta}$, the error between $V^*_{n,\delta}(\sigma_{n,\delta})$ and $\hat{V}^*_{n,\delta}(\sigma_{n,\delta})$ satisfies*

$$0 \le V^*_{n,\delta}(\sigma_{n,\delta}) - \hat{V}^*_{n,\delta}(\sigma_{n,\delta}) \le N_q\overline{\lambda}h_y^{(1)}\delta^y + \frac{\epsilon}{N}$$

given that the discretized observations w are chosen so that

$$\gamma_\delta(w|z') > \gamma(w|z')|\Theta(w)|,$$

and with $\overline{\lambda}$ the largest Lebesgue measure of sets $\overline{K}_{y^q}$.

Lemma 10 requires defining the representative points $w = (y^{x,i,y^q}, y^q)$ so that the integral of γ over $\Theta(w)$ is greater than a piecewise constant approximation integrated over $\Theta(w)$, which can be satisfied by picking the points y^{x,i,y^q} where the Gaussian density γ_x is minimized within cell $\overline{K}_{i,y^q}$. Without this requirement, finding $\alpha^{*(w)}_{n+1,\delta}$ at a finite set of points still guarantees a lower bound to the value function for any time n, and is intuitively more accurate as $\delta^y \to 0$.

Lemma 10 leads to the following theorem regarding the error between $V^*_{n,\delta}(\sigma_{n,\delta})$ (based on continuous observations) and $\tilde{V}^*_{n,\delta}(\sigma_{n,\delta})$ (based on discretized observations). We again use the notation $\tilde{V}$ to indicate that $\tilde{V}$ is represented by the set $\tilde{\Gamma}$ of α-functions calculated using the discretized observations.

Theorem 6. *Given discretized observation process W_δ with transition function* (43)*, for any time $n \in [0, N]$, the error between $V^*_{n,\delta}(\sigma_{n,\delta})$ calculated according to $\mathcal{Y}$ and $\tilde{V}^*_{n,\delta}(\sigma_{n,\delta})$ calculated according to W_δ satisfies*

$$0 \le V^*_{n,\delta}(\sigma_{n,\delta}) - \tilde{V}^*_{n,\delta}(\sigma_{n,\delta}) \le (N-n)N_q\overline{\lambda}h_y^{(1)}\delta^y + \frac{(N-n)\epsilon}{N}$$

for any $\sigma_{n,\delta} \in \Sigma_{n,\delta}$, with $\overline{\lambda}$ the largest Lebesgue measure of sets $\overline{K}_{y^q}$.

Specifically, the safety probability for $\hat{\mathcal{H}}$ satisfies

$$0 \leq p^N_{\text{safe},\delta}(\rho_\delta; K_\delta) - \tilde{V}^*_{0,\delta}(\rho_\delta) \leq N N_q \overline{\lambda} h^{(1)}_y \delta^y + \epsilon.$$

2) Discretized observations for Gaussian mixture approximation: The results of discretizing the observations for the Gaussian mixture abstraction are nearly identical to those for the finite state abstraction. The main difference arises in approximating the integral $Y(\Theta(w)|s')$ with a Gaussian sum. To ensure the approximate value function provides a lower bound to $V^*_{n,g}$, we must under-approximate the integral $Y(\Theta(w))$ for each w. We define $\tilde{V}^*_{n,g}$ similarly to $\tilde{V}^*_{n,\delta}$, using discretized observations and the RBF approximation to the indicator function,

$$\begin{aligned}
\tilde{\alpha}_{n,g}(s) &= \sum_{w \in W_\delta} \int_{\mathcal{S}} \tilde{\alpha}^{*(w)}_{n+1,g}(s') \gamma_g(w|s') \tau(s'|s,u)\, ds' \sum_{i=1}^{I_q} w_i(q) \phi_i(x) \\
\tilde{\alpha}^{*(w)}_{n+1,g}(s') &= \arg \left\{ \sup_{\tilde{\alpha}_{n+1,g} \in \tilde{\Gamma}_{n+1,g}} \sum_{q \in \mathcal{Q}} \int_{\mathbb{R}^m} \tilde{\alpha}^{*(w)}_{n+1,g}(s') \gamma_g(w|s') \tau(s'|s,u)\, ds' \left[\sum_{i=1}^{I_q} w_i(q) \phi_i(x) \right] \sigma_{n,g}(x,q)\, dx \right\}
\end{aligned} \tag{46}$$

and $\hat{V}^*_{n,g}$ is the intermediate value function that finds the optimal $\alpha^{*(w)}_{n+1,g} \in \Gamma_{n+1,g}$, rather than in $\tilde{\Gamma}_{n+1,g}$. We can bound the error between $V^*_{n,g}$ and $\hat{V}^*_{n,g}$, and between $V^*_{n,g}$ and $\tilde{V}^*_{n,g}$, equivalently to Lemma 10 and Theorem 6, respectively.

Lemma 11. *For any time $n \in [0, N]$ and $\sigma_{n,g} \in \Sigma_{n,g}$, the error between $V^*_{n,g}(\sigma_{n,g})$ and $\hat{V}^*_{n,g}(\sigma_{n,g})$ satisfies*

$$0 \leq V^*_{n,g}(\sigma_{n,g}) - \hat{V}^*_{n,g}(\sigma_{n,g}) \leq N_q \overline{\lambda} h^{(1)}_y \overline{\alpha}_{n+1} \delta^y + \frac{\epsilon}{N}$$

given that the observations w are chosen so that

$$\gamma_g(w|s') > \gamma(w|s')|\Theta(w)|,$$

and with $\overline{\lambda}$ the largest Lebesgue measure of sets $\overline{K}_{y^q}$.

Theorem 7. *Given discretized observation process W_δ with transition function* (44)*, for any time $n \in [0, N]$, the error between $V^*_{n,g}(\sigma_{n,g})$ calculated according to $\mathcal{Y}$ and $\tilde{V}^*_{n,g}(\sigma_{n,g})$ calculated according to W_δ satisfies*

$$0 \leq V^*_{n,g}(\sigma_{n,g}) - \tilde{V}^*_{n,g}(\sigma_{n,g}) \leq N_q \overline{\lambda} h^{(1)}_y \left(\sum_{i=n+1}^{N} \overline{\alpha}_i \right) \delta^y + \frac{(N-n)\epsilon}{N}$$

for any $\sigma_{n,g} \in \Sigma_{n,g}$, with $\overline{\lambda}$ the largest Lebesgue measure of sets $\overline{K}_{y^q}$.

Specifically, the safety probability for the Gaussian mixture approximation satisfies

$$0 \leq p^N_{\text{safe},g}(\rho; K) - \tilde{V}^*_{0,g}(\rho) \leq N_q \overline{\lambda} h^{(1)}_y \left(\sum_{i=1}^{N} \overline{\alpha}_i \right) \delta^y + \epsilon.$$

The proofs of Lemma 11 and Theorem 7 follow directly from the proofs of Lemma 10 and Theorem 6.

To summarize, given either the finite state or Gaussian mixture approximation, we can subsequently 1) sample y from $\mathcal{Y}$ and u from $\mathcal{U}$ to generate the progressive subsets $\tilde{\Sigma}_{n,\delta}$ or $\tilde{\Sigma}_{n,g}$, and 2) discretize $\mathcal{Y}$ and use the set W_δ to calculate $\tilde{\alpha}^\delta_{w,u,\sigma_\delta}$ or $\tilde{\alpha}^g_{w,u,\sigma_g}$, which are then used to generate $\tilde{\alpha}_{n,\delta} \in \tilde{\Gamma}_{n,\delta}$ and $\tilde{\alpha}_{n,g} \in \tilde{\Gamma}_{n,g}$. Using sets $\tilde{\Sigma}_{n,\delta}$ and

$\tilde{\Gamma}_{n,\delta}$ in place of $\Sigma_{n,\delta}$ and $\Gamma_{n,\delta}$ provides a lower bound to the safety probability $p^N_{\text{safe},\delta}(\rho_\delta; K_\delta)$ that converges as δ^σ and δ^y approach zero (and similarly for $\tilde{\Sigma}_{n,g}$ and $\tilde{\Gamma}_{n,g}$).

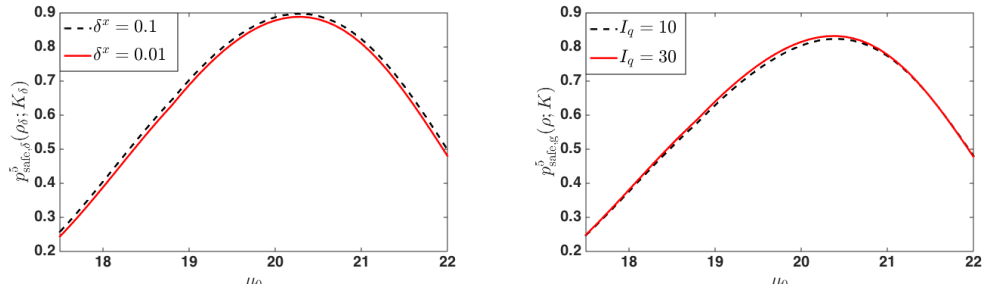

(a) (b)

Fig. 1: Comparison of safety probabilities over varying initial distribution $\rho = \phi(x; \mu_0, 1)$ and $q_0 = 0$, using the finite state approximation (a) and Gaussian mixture approximation (b). In both (a) and (b) $\delta^y = 0.5$. Fig. (a) compares probabilities for $\delta^x = 0.1$ (black dashed line) and $\delta^x = 0.01$ (red solid line). Fig. (b) compares probabilities for $I_q = 10$ (black dashed line) and $I_q = 30$ (red solid line). The refinement of δ^x and increase in I_q have a small impact on the safety probabilities.

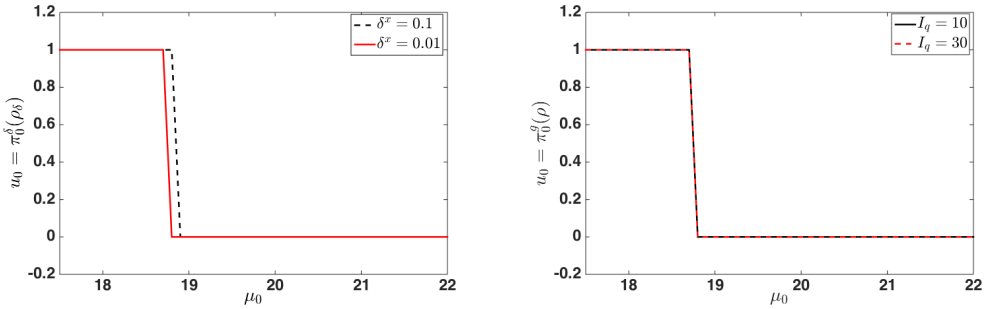

(a) (b)

Fig. 2: Comparison of optimal control inputs as a function of $\rho = \phi(x; \mu_0, 1)$ with $q_0 = 0$, using the finite state approximation (a) and Gaussian mixture approximation (b). In both (a) and (b), $\delta^y = 0.5$. Fig. (a) compares control inputs for $\delta^x = 0.1$ (black dashed line) and $\delta^x = 0.01$ (red solid line). Fig. (b) compares control inputs for $I_q = 10$ (black dashed line) and $I_q = 30$ (red solid line), which in this case are the same. All approaches produce a thresh-hold policy that turns the heater off for $\mu_0 > \approx 18.7$.

V. Numerical Example

The temperature regulation problem is a benchmark example for hybrid systems, and a stochastic version with perfect state information is presented in [4]. We consider the case of one heater, which can either be turned on to heat one room, or turned off. The temperature of the room at time n is given by the continuous variable x_n, and the discrete state $q_n = 1$ indicates the heater is on at time n, and $q_n = 0$ denotes the heater is off. The stochastic difference equation governing the temperature is given by

$$x_{n+1} = (1 - b)x_n + cq_{n+1} + bx_a + v_n$$

with constants $b = 0.0167$, $c = 0.8$, and $x_a = 6$, and v_n i.i.d. Gaussian random variables with mean zero and variance v^2. The control input is given by $u_n \in \mathcal{U}$ with $\mathcal{U} = \{0, 1\}$, but the chosen control is not always implemented with probability 1. Instead, q_n is updated probabilistically, dependent on u_{n-1} and q_{n-1}, with transition function $T_q(q_{n+1}|q_n, u_n)$. So while function $\pi_n(\sigma_n)$ deterministically returns a single control input, control input $u_n = \pi_n(\sigma_n)$ may not always be implemented.

To model this as a partially observable problem, assume the actual temperature is unknown, and only a noisy measurement is available to the controller. The controller does, however, know whether the heater is on or off at time n (i.e. q_n is perfectly observed). The observation $y_n = y_n^x$ is given by $y_n^x = x_n + w_n$, with w_n i.i.d. Gaussian random variables with mean zero and variance w^2.

It is desirable to keep the temperature of the room between 17.5 and 22 degrees Celsius at all times, hence the safe region $K = [17.5, 22]$ does not depend on the discrete state q_n (so $\mathbf{1}_K(s) = \mathbf{1}_K(x)$). We consider the probability of remaining within K for $N = 5$ time steps given initial temperature distribution ρ_x normally distributed with varying mean $\mu_0 \in K$ and variance $\mathcal{P}_0 = 1$. The initial mode is given as $q_0 = 0$. The finite state and Gaussian mixture approximations are used in a PBVI algorithm in the style of Perseus [22].

We consider a uniform grid ($\delta_{i,q}^x = \delta^x$ constant for all i, q) over the region $K \subset \mathbb{R}$ for the finite state approximation, with representative points at the end-point of each grid cell. For example, setting $\delta^x = 0.1$ gives $x^{1,q} = 17.5, x^{2,q} = 17.6, \ldots$ for $q = 0$ and $q = 1$, and a total of $m_q = 45$ cells $K_{i,q}$. The function $\xi(x, q)$ maps q to itself, and maps x to the nearest $x^{i,q}$ in absolute value.

The Gaussian mixture approximation utilizes an RBF approximation of the indicator function calculated using MATLAB's *gmdistribution* function. We used a reduction process to limit the number of components of each α and σ for the Gaussian mixture approximation. Similar Gaussians are combined into a single component based on the 2-norm distance between functions [29]. Each mixture was limited to 30 components to reduce overall computation time without overly sacrificing accuracy. This number can easily be changed, however, depending on the importance of speed versus accuracy.

Both approximations employ a sequence of sampled sets $\tilde{\Sigma}_n$ and a finite set of observations to calculate the α-functions for the PBVI algorithm. To generate the sets $\tilde{\Sigma}_n$, we initialized a set of 40 states σ_0 normally distributed with variance $\mathcal{P}_0$ and mean μ_0 randomly chosen uniformly on K. Each σ_0 was updated according to $\Phi_{y,u}^g$ or $\Phi_{y,u}^\delta$ with u chosen randomly and y sampled from the corresponding σ_0 (i.e. $y \sim p(y|\sigma_0, u)$). This process was repeated N times, so that for each time step we had a set of 40 sampled σs. The finite set of observations were produced by a uniform grid over $\overline{K} = [16, 24]$, again using end-points as the representative observations.

To compare performance of the finite state and Gaussian mixture approximations, we present computation times and safety probability estimates for each, with varying δ^x, δ^y, and number of components in the indicator approximation. Safety probabilities for varying initial distributions ρ are presented in Figs. 1a and 1b for the finite state approximation and Gaussian mixture approximation, respectively. The optimal policy at time zero is shown for varying ρ in Figs. 2a and 2b for the finite and Gaussian approximations, respectively. Computation times for the Gaussian mixture approximation are given in Table I, and for the finite state approximation in Table II.

We also show sample RBF approximations to the indicator function $\mathbf{1}_K$ in Fig. 3 with varying numbers of

components I_q. The error between the RBF approximation and $\mathbf{1}_K$ for varying I_q is shown in Fig. 4. As the number of components increases, the approximation becomes more accurate, although as seen in Fig. 3, oscillations remain at the endpoints of K. The increasing accuracy is most apparent in Fig. 4, and demonstrates the convergence towards zero of the error δ^I with increasing I_q.

We show safety probabilities for $\delta^y = 0.5$. Decreasing δ^y causes a slight increase in the safety probabilities, as expected, but there is not a significant improvement in the probability estimates, although as seen in Tables I and II, the increase in computation time is significant. This is likely problem-specific, and the value of δ^y may have a greater impact for some applications.

The safety probability estimates for the finite state approximation are in general greater than for the Gaussian mixture approximation. The mixture reduction method employed, as well as the indicator function approximation, make the Gaussian method seemingly less accurate than the finite state approximation. However, over a finer mesh δ^x, the finite state method results in greater computation time. Note also that for $\delta^x = 0.01$, the safety probabilities decrease relative to the probabilities for $\delta^x = 0.1$, which highlights that although sampling from the information space and discretizing the observations guarantees a lower bound, no such guarantee exists when discretizing the state space, and there is a chance the safety estimates do not bound the true safety probabilities from below.

Although the coarse grid produces similar results to the fine grid ($\delta^x = 0.1$ versus $\delta^x = 0.01$), in higher dimensional problems the number of grid cells becomes prohibitive even when δ^x is large, and the Gaussian mixture approximation may be more computationally tractable. All scenarios produce a nearly identical optimal thresh-hold policy based on the initial mean μ_0, indicating that an optimal policy may be computed fairly quickly using any of the above methods.

The computation time is unfortunately still quite high in both cases, and at this time each method is likely applicable to systems with only a few modes and continuous state dimensions of no more than two or three. It is likely, however, that computation time can be improved by using more sophisticated point-based solvers, such as [30]. Further gains may be possible through adaptive gridding techniques, similar to those in [13], to decrease the number of finite states required without sacrificing accuracy.

Interestingly, increasing the number of components in the RBF approximation to the indicator function only slightly improves the safety estimates of the Gaussian mixture approximation, although the error from increasing the number of components to 30 drops significantly. This may be caused by the mixture reduction technique, leading to a loss in the added benefit of an increased number of components when that number is again reduced.

	$I_q = 10$		$I_q = 30$	
	$\delta^y = 1.0$	$\delta^y = 0.5$	$\delta^y = 1.0$	$\delta^y = 0.5$
Comp. time (s)	365.5	1625.9	1865.0	5586.1

TABLE I: Computation times using PBVI with Gaussian mixture approximation, for varying number of components I_q for RBF approximation to $\mathbf{1}_K$ and discretized observation spacing δ^y.

	$\delta^x = 0.1$			$\delta^x = 0.01$		
	$\delta^y = 1.0$	$\delta^y = 0.5$	$\delta^y = 0.1$	$\delta^y = 1.0$	$\delta^y = 0.5$	$\delta^y = 0.1$
Comp. time (s)	50.5	205.1	1599.8	8961.1	15343.7	108591.3

TABLE II: Computation times using PBVI with finite state approximation, for varying continuous state spacing δ^x and discretized observation spacing δ^y.

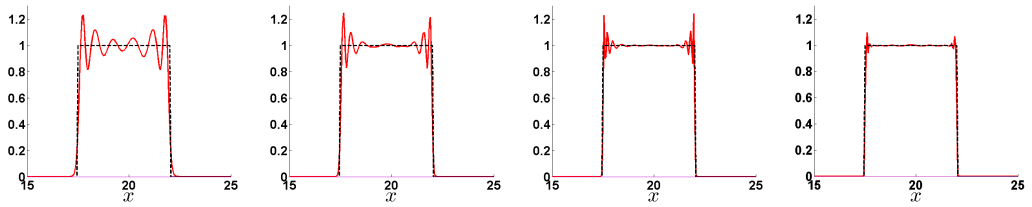

(a) (b) (c) (d)

Fig. 3: Comparison between $\mathbf{1}_K(x)$ (in black, dashed line) to RBF approximation (red, solid line) for (a) $I_q = 10$ components, (b) $I_q = 30$ components, (c) $I_q = 100$ components, and (d) $I_q = 400$ components. As the number of components increases, the approximation improves, although oscillations at the endpoints remain.

However, although the error with $I_q = 10$ is large, we obtain safety estimates that are quite similar to the finite state approximation. This requires further investigation, but may help in decreasing computation time without losing significant accuracy by choosing I_q to be small.

VI. Conclusion

We have presented the first numerical results for verification and controller synthesis for safety objectives, given a partially observable DTSHS. We have considered two approximations that enable the use of a well-known POMDP optimization technique. The first approximation discretizes the state space over a compact set K and enables a vector representation of the information states and α-functions. The second approximates the indicator function over

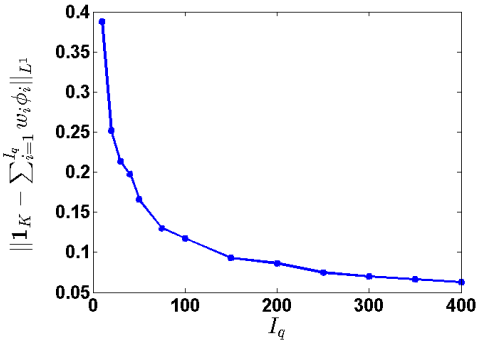

Fig. 4: The integrated error for RBF approximations to indicator function $\mathbf{1}_K$ with a varying number of components I_q.

compact set K using a finite set of Gaussian radial basis functions and enables a Gaussian mixture representation of the information states and α-functions. We can apply point-based value iteration to either approximation, and guarantee a lower bound to the safety probability, which is proven to converge to the true safety probability of the original PODTSHS. A simple numerical example shows that both methods provide similar safety estimates. The finite state approximation is faster when a coarse discretization is used, but quickly becomes slower than the Gaussian mixture approximation with a finer discretization. Therefore, although the Gaussian mixture produces lower safety estimates, it may be better suited to higher dimensional problems.

Although we present a switched affine system with additive Gaussian noise, both approximations may be extended to non-Gaussian systems. Convergence results for the finite state approximation apply to arbitrary transition kernels T_x and Y_x, given they still satisfy certain Lipschitz conditions. The Gaussian mixture approximation further requires approximating T_x, T_q, and Y_x with Gaussian mixtures, and introduces additional error. We also focus on the safety problem, although the computational techniques presented will apply to other verification properties such as reachability, reach-avoid objectives, and others, by modifying the information state and α-functions slightly. We are currently working to formally extend these results to other verification objectives and more complex applications. Because both methods are relatively slow, we plan to continue to refine them to decrease computation time, which is possible through the use of more sophisticated existing point-based solvers. We are also exploring more efficient computation by exploiting problem structure, through the use of adaptive gridding schemes and other representations of α and σ beyond vectors and Gaussian mixtures.

APPENDIX

Proof of Lemma 1: By induction. At time N,

$$V_N^*(\sigma_N) = \int_{\mathcal{S}} \mathbf{1}_K(s)\sigma_N(s)\,ds.$$

By defining $\alpha_N(s) = \mathbf{1}_K(s)$, we obtain the desired result. Note that this definition of α_N is in line with the definition given in Section II-B, because although it does not represent a full policy tree (being at the terminal time, there are no more branches on the tree), it does represent the immediate value of being in state (x, q), given by $\mathbf{1}_K(x, q)$.

Next, assuming $V_{n+1}^*(\sigma_{n+1}) = \sup_{\alpha_{n+1}} \langle \alpha_{n+1}, \sigma_{n+1} \rangle$, V_n^* can be written as

$$\begin{aligned} V_n^*(\sigma_n) &= \max_{u\in\mathcal{U}} \mathbb{E}^\pi \left[\sup_{\alpha_{n+1}\in\Gamma_{n+1}} \langle \alpha_{n+1}, \sigma_n \rangle \right] \\ &= \max_{u\in\mathcal{U}} \int_{\mathcal{Y}} \sup_{\alpha_{n+1}\in\Gamma_{n+1}} \langle \alpha_{n+1}, \Phi_{y,u}\sigma_n \rangle \mathbb{P}(dy|\sigma_n, u) \\ &= \max_{u\in\mathcal{U}} \int_{\mathcal{Y}} \sup_{\alpha_{n+1}\in\Gamma_{n+1}} \left[\int_K \int_{\mathcal{S}} \alpha_{n+1}(s')\gamma(y|s')\tau(s'|s,u)\sigma_n(s)\,ds'\,ds \right] dy. \end{aligned}$$

Then for a specific observation y, action u, and function α_{n+1}^i (index i used to identify a particular $\alpha_{n+1} \in \Gamma_{n+1}$), the function $\alpha_{y,u}^i$ can be defined as

$$\alpha_{y,u}^i(s) = \int_{\mathcal{S}} \alpha_{n+1}^i(s')\gamma(y|s')\tau(s'|s,u)\,ds'\mathbf{1}_K(s). \tag{47}$$

Because $\alpha^i_{y,u}$ does not depend on σ_n, we can redefine the supremum over all Γ_{n+1} to be over all $\alpha^i_{y,u}$.

$$V^*_n(\sigma_n) = \max_{u\in\mathcal{U}} \int_{\mathcal{Y}} \sup_{\{\alpha^i_{y,u}\}} \langle \alpha^i_{y,u}, \sigma_n\rangle \, dy$$

For a specific σ, u, and y, we define

$$\begin{aligned} \alpha_{y,u,\sigma}(s) &= \arg\sup_i \langle \alpha^i_{y,u}, \sigma\rangle \\ &= \int_{\mathcal{S}} \alpha^{*(y)}_{n+1}(s')\gamma(y|s')\tau(s'|s,u)\, ds' \mathbf{1}_K(s) \end{aligned} \tag{48}$$

with $*(y)$ denoting the index i of the α-function $\alpha^i_{y,u}$ that maximizes the inner product. We further simplify V^*_n as

$$\begin{aligned} V^*_n(\sigma_n) &= \max_{u\in\mathcal{U}} \int_{\mathcal{Y}} \langle \alpha_{y,u,\sigma}, \sigma_n\rangle \, dy \\ &= \max_{u\in\mathcal{U}} \left\langle \int_{\mathcal{Y}} \alpha_{y,u,\sigma}\, dy, \sigma_n \right\rangle. \end{aligned} \tag{49}$$

Therefore, the set of all $\{\alpha_n\}$ is described by

$$\Gamma_n = \bigcup_{\sigma} \left\{ \int_{\mathcal{Y}} \alpha_{y,u^*,\sigma}\, dy \right\} \tag{50}$$

with u^* the control inputs chosen according to (49), and V^*_n may be written as

$$V^*_n(\sigma_n) = \sup_{\alpha_n\in\Gamma_n} \langle \alpha_n, \sigma_n\rangle \tag{51}$$

■

Proof of Theorem 1: By induction. At time $n = 0$, $\sigma_0(s) = \rho(s) = \hat{\sigma}_0(s)$ and the inequality is trivially satisfied. For all $i = 0, \ldots, n$, assume $|\sigma_i(s) - \hat{\sigma}_i(s)| \le \eta^\sigma_i \delta^x$. At time $i = n+1$, for any $y \in \mathcal{Y}$ and any $u \in \mathcal{U}$,

$$|\sigma_{n+1}(s') - \hat{\sigma}_{n+1}(s')| \le \min\left\{ \frac{1}{p(y|\sigma_n,u)}, \frac{1}{p(y|\hat{\sigma}_n,u)} \right\} \Bigg| \gamma(y|s') \int_K \tau(s'|s,u)\sigma_n(s)\, ds$$

$$- \gamma(y|\xi(s')) \int_K \tau(s'|\xi(s),u)\hat{\sigma}_n(s)\, ds \Bigg|.$$

We add and subtract $\gamma(y|\xi(s')) \int_K \tau(s'|s,u)\sigma_n(s)\, ds$ and $\gamma(y|\xi(s')) \int_K \tau(s'|\xi(s),u)\sigma_n(s)\, ds$, apply the triangle inequality, and use the Lipschitz inequalities (11).

$$\begin{aligned} |\sigma_{n+1}(s') - \hat{\sigma}_{n+1}(s')| &\le \min\left\{ \frac{1}{p(y|\sigma_n,u)}, \frac{1}{p(y|\hat{\sigma}_n,u)} \right\} \Bigg[|\gamma(y|s') - \gamma(y|\xi(s'))| \int_K \tau(s'|s,u)\sigma_n(s)\, ds \\ &\quad + \gamma(y|\xi(s')) \int_K |\tau(s'|s,u) - \tau(s'|\xi(s),u)|\, \sigma_n(s)\, ds + \gamma(y|\xi(s')) \int_K \tau(s'|\xi(s),u)\, |\sigma_n(s) - \hat{\sigma}_n(s)|\, ds \Bigg] \\ &\le \min\left\{ \frac{1}{p(y|\sigma_n,u)}, \frac{1}{p(y|\hat{\sigma}_n,u)} \right\} \Bigg[h^{(2)}_y \|x' - \xi(x')\| \phi^*_v + \phi^*_w h^{(2)}_x \|x' - \xi(x')\| + \phi^*_w h_q \|x - \xi(x)\| \phi^*_v \\ &\quad + \phi^*_w |\sigma_n - \hat{\sigma}_n| \sum_{q\in\mathcal{Q}} T_q(q'|s,u) \int_{K_q} \tau_x(x'|q',\xi(s),u)\, dx \Bigg] \end{aligned}$$

Since τ_x is bounded by ϕ^*_v, and the Lebesgue measure of K_q is at most λ, we obtain

$$|\sigma_{n+1}(s') - \hat{\sigma}_{n+1}(s')| \le \min\left\{ \frac{1}{p(y|\sigma_n,u)}, \frac{1}{p(y|\hat{\sigma}_n,u)} \right\} \left[\phi^*_v h^{(2)}_y \delta^x + \phi^*_w h^{(2)}_x \delta^x + \phi^*_w \phi^*_v h_q \delta^x + \phi^*_w N_q \lambda |\sigma_n - \hat{\sigma}_n| \right].$$

Combining terms and applying the induction hypothesis gives the desired result. ■

Proof of Lemma 3: We exploit properties of the derivative of a Gaussian distribution, which bounds the Lipschitz constants for γ_x from above. For clarity of notation, we will write C rather than $C(q)$, with dependence on q implicit. The constant $h_y^{(2)}$ is the maximum value of the derivative of $\phi(y^x; Cx, \mathcal{W})$ with respect to x:

$$\left\| \frac{\partial \phi}{\partial x} \right\| \leq \frac{1}{(2\pi)^{\frac{n}{2}} |\mathcal{W}|^{\frac{1}{2}} \lambda_1^w} \|y^x - Cx\| \|C\| e^{-\frac{\|y^x - Cx\|^2}{2\lambda_1^w}}. \tag{52}$$

Recall that $x^T A^{-1} x \geq \lambda_{\min}(A^{-1}) \|x\|^2 = \frac{\|x\|^2}{\lambda_{\max}(A)}$ for A a symmetric matrix. Therefore $\exp(-\frac{1}{2}(y-Cx)^T \mathcal{W}^{-1}(y-Cx)) \leq \exp(-\frac{\|y-Cx\|^2}{2\lambda_{\max}(\mathcal{W})})$. Taking the derivative again with respect to $\|y - Cx\|$ and setting equal to zero, we see that the maximum of (52) occurs at $\|y - Cx\| = \sqrt{\lambda_1^w}$.

Although $\|\frac{\partial \phi}{\partial x}\| \leq h_y^{(2)}$, we create a tighter bound for the case in which $\|y^x - Cx\|$ is greater than $\sqrt{\lambda_1^w}$ (for $y^x \in \mathbb{R}^l$ such that there exists $x \in K_{i,q}$ for which $\|y^x - Cx\| = \sqrt{\lambda_1^w}$, the upper bound $h_y^{(2)}$ is attained) using the following function.

$$h_y(y^x) = \max_{x \in K_{i,q}} \left\{ \frac{\|C\|}{(2\pi)^{\frac{n}{2}} |\mathcal{W}|^{\frac{1}{2}} \lambda_1^w} \|y^x - Cx\| e^{-\frac{\|y^x - Cx\|^2}{2\lambda_1^w}} \right\}.$$

Then,

$$\begin{aligned}
\int_{\mathbb{R}^l} |\gamma_x(y^x|x,q) - \gamma_x(y^x|\overline{x},q)| \, dy^x &\leq \int_{\mathbb{R}^l} h_y(y^x) \|x - \overline{x}\| \, dy^x \\
&\leq \delta_{i,q}^x \int_{\{y^x : \|y^x - Cx\|^2 \leq \lambda_1^w, y^x \in \mathbb{R}^l, x \in K_{i,q}\}} h_y^{(2)} \, dy^x \\
&\quad + \delta_{i,q}^x \int_{\{y^x : \|y^x - Cx\|^2 > \lambda_1^w, y^x \in \mathbb{R}^l, x \in K_{i,q}\}} h_y(y^x) \, dy^x \\
&= \delta_{i,q}^x \beta_{1,i,q}^y h_y^{(2)} + \delta_{i,q}^x \int_{\{y^x : \|y^x - Cx\|^2 > \lambda_1^w, y^x \in \mathbb{R}^l, x \in K_{i,q}\}} h_y(y^x) \, dy^x
\end{aligned} \tag{53}$$

We use the change of variable $v = \|y^x - Cx^*\|$, with $x^* = \arg\min_{x \in K_{i,q}} \|y^x - Cx\|$, to rewrite the second term of (53), and apply an identity for integrals of polynomials.

$$\begin{aligned}
&\int_{\{y^x : \|y^x - Cx\|^2 > \lambda_1^w, y^x \in \overline{K}, x \in K_{i,q}\}} h_y(y^x) \, dy^x \\
&\quad = \frac{\|C\|}{(2\pi)^{\frac{n}{2}} |\mathcal{W}|^{\frac{1}{2}} \lambda_1^w} \int_{\sqrt{\lambda_1^w}}^{\infty} v e^{-\frac{v^2}{2\lambda_1^w}} \, dv \\
&\quad \leq \frac{\|C\|}{(2\pi)^{\frac{n}{2}} |\mathcal{W}|^{\frac{1}{2}} \lambda_1^w} \int_{0}^{\infty} v e^{-\frac{v^2}{2\lambda_1^w}} \, dv \\
&\quad \leq \frac{\|C\|}{(2\pi)^{\frac{n}{2}} |\mathcal{W}|^{\frac{1}{2}}}
\end{aligned} \tag{54}$$

Inserting (54) into (53) proves the lemma. ■

Proof of Lemma 5: By induction. At time N,

$$|\alpha_N^i(s) - \tilde{\alpha}_N^i(s)| = \left| \int_{\mathcal{S}} (\mathbf{1}_K(s) - \mathbf{1}_{K_\delta}(\xi(s))) \, ds \right| = 0$$

since for any $s \in K$, by definition $\xi(s) \in K_\delta$. Assume for all $j = N-1, \ldots, n+1$, $|\alpha_j^i(s) - \tilde{\alpha}_j^i(s)| \leq (N-n)N_q \left[\beta_1^y h_y^{(2)} + \beta_1^x h_x^{(2)} + \beta_2^y + \beta_2^x + h_q\right] \delta^x$. For $j = n$, separate terms and apply the triangle inequality as in the proof of Theorem 1.

$$\begin{aligned}
|\alpha_n^i(s) - \tilde{\alpha}_n^i(s)| = \Bigg| \int_{\mathcal{S}} \int_{\mathcal{Y}} & \alpha_{n+1}^{i(y)}(s') \gamma(y|s') \tau(s'|s, u^i) \mathbf{1}_K(s) \, dy \, ds' \\
& - \int_{\mathcal{S}} \int_{\mathcal{Y}} \tilde{\alpha}_{n+1}^{i(y)}(s') \gamma(y|\xi(s')) \tau(s'|\xi(s), u^i) \mathbf{1}_{K_\delta}(\xi(s)) \, dy \, ds' \Bigg| \\
\leq & \left| \alpha_{n+1}^{i(y)}(s') - \tilde{\alpha}_{n+1}^{i(y)}(s') \right| + N_q \left[\beta_1^y h_y^{(2)} + \beta_2^y\right] \delta^x + N_q \left[h_q + \beta_1^x h_x^{(2)} + \beta_2^x\right] \delta^x
\end{aligned} \tag{55}$$

The second term of (55) comes from Lemma 3 and noting that $\alpha(s)$ represents a probability that is bounded above by one. The third term comes from Lemma 4 and the Lipschitz inequality for T_q (11). The term $\mathbf{1}_K(s)$ does not affect the bound, and only indicates that both $\alpha_n(s)$ and $\tilde{\alpha}_n(s)$ are equal to zero for $s \notin K$. Applying the induction hypothesis to (55) gives the desired result. ■

Proof of Theorem 2: By construction. At any time $n \in [0, N]$, given $\sigma_n \in \Sigma$ and $\sigma_{n,\delta} \in \Sigma_\delta$, we can rewrite the value function evaluated at σ in terms of α-functions.

$$\begin{aligned}
\left| V_n^*(\sigma_n) - V_{n,\delta}^*(\sigma_{n,\delta}) \right| &= \left| \sup_{\alpha_n \in \Gamma_n} \langle \alpha_n, \sigma_n \rangle - \sup_{\alpha_{n,\delta} \in \Gamma_{n,\delta}} \langle \alpha_{n,\delta}, \sigma_{n,\delta} \rangle \right| \\
&= \left| \langle \alpha_n^k, \sigma_n \rangle - \langle \alpha_{n,\delta}^l, \sigma_{n,\delta} \rangle \right|
\end{aligned}$$

Assume without loss of generality that $\langle \alpha_n^k, \sigma_n \rangle \geq \langle \alpha_{n,\delta}^l, \sigma_{n,\delta} \rangle$. Then, because $\langle \tilde{\alpha}_n^k, \hat{\sigma}_n \rangle \leq \langle \hat{\alpha}_n^l, \hat{\sigma}_n \rangle$ by definition of the optimality of $\hat{\alpha}_n^l$, we can write

$$\begin{aligned}
\left| V_n^*(\sigma_n) - V_{n,\delta}^*(\sigma_{n,\delta}) \right| &= \langle \alpha_n^k, \sigma_n \rangle - \langle \alpha_{n,\delta}^l, \sigma_{n,\delta} \rangle \\
&\leq \langle \alpha_n^k, \sigma_n \rangle - \langle \tilde{\alpha}_{n,\delta}^k, \sigma_{n,\delta} \rangle \\
&\leq \left| \langle \alpha_n^k, \sigma_n \rangle - \langle \tilde{\alpha}_n^k, \sigma_n \rangle \right| + \left| \langle \tilde{\alpha}_n^k, \sigma_n \rangle - \langle \tilde{\alpha}_n^k, \hat{\sigma}_n \rangle \right| \\
&\leq \int_{\mathcal{S}} \left| \alpha_n^k(s) - \tilde{\alpha}_n^k(s) \right| \sigma_n(s) \, ds + \int_{\mathcal{S}} \tilde{\alpha}_n^k(s) \left| \sigma_n(s) - \hat{\sigma}_n(s) \right| \, ds
\end{aligned} \tag{56}$$

Applying Lemma 5 to the first term of (56), and noting that the integral in the second term is in fact taken over K rather than $\mathcal{S}$ since $\tilde{\alpha}_n^k(s)$ is zero for all $s \notin K$, we obtain

$$\left| V_n^*(\sigma_n) - V_{n,\delta}^*(\sigma_{n,\delta}) \right| \leq (N-n) N_q \left[h_q + \beta_1^y h_y^{(2)} + \beta_1^x h_x^{(2)} + \beta_2^y + \beta_2^x\right] \delta^x + N_q \lambda \eta_n^\sigma \delta^x \tag{57}$$

which completes the proof. ■

Proof of Lemma 6: The proof of both Lemma 6 and Lemma 7 require the following Gaussian density identities. First, the product of two Gaussian densities is again a Gaussian density, up to a constant factor.

$$\phi(x; \mu_1, \mathcal{P}_1) \phi(x; \mu_2, \mathcal{P}_2) = \phi(\mu_1; \mu_2, \mathcal{P}_1 + \mathcal{P}_2) \phi(x; \tilde{\mu}, \tilde{\mathcal{P}}) \tag{58}$$

with

$$\begin{aligned}
\tilde{\mu} &= \tilde{\mathcal{P}} \left(\mathcal{P}_1^{-1} \mu_1 + \mathcal{P}_2^{-1} \mu_2 \right) \\
\tilde{\mathcal{P}} &= \left(\mathcal{P}_1^{-1} + \mathcal{P}_2^{-1} \right)^{-1}
\end{aligned} \tag{59}$$

Second, for invertible matrix A, constant b, and variables x and y,

$$\phi(y; Ax + b, \mathcal{P}) = |A^{-1}|\phi(x; A^{-1}(y - b), A^{-1}\mathcal{P}A^{-T}). \tag{60}$$

Both identities are easily shown directly.

The proof then follows by construction. Given $\sigma_{n,g}(x, q) = \sum_{l=1}^{L} w_{l,n}^{\sigma}(q)\phi(x; \mu_{l,n}^{\sigma}(q), \mathcal{P}_{l,n}^{\sigma}(q))$, observation $y \in \mathcal{Y}$, and control input $u \in \mathcal{U}$, the operator $\Phi_{y,u}^{g}\sigma_{n,g}$ produces

$$\Phi_{y,u}^{g}\sigma_{n,g} = \gamma(y|s', u) \sum_{q\in\mathcal{Q}} \int_{\mathbb{R}^m} \left[\sum_{i=1}^{I_q} w_i^I(q)\phi_i^I(x)\right] \tau(s'|s, u)\sigma_{n,g}(s)\, ds. \tag{61}$$

We use the notation w_i^I, μ_i^I, and $\mathcal{P}_i^I$ for the weights, means, and covariances of the indicator function approximation to distinguish them from the other Gaussian mixtures representing σ and α. Replacing $\sigma_{n,g}$ by its Gaussian mixture representation in (61), and expanding γ and τ, gives

$$\begin{aligned}
\left(\Phi_{y,u}^{g}\sigma_{n,g}\right)(Cx', q') &= Y_q(y^q|q')\phi(y^x; Cx', \mathcal{W}) \sum_{q\in\mathcal{Q}} \int_{\mathbb{R}^m} \left[\sum_{i=1}^{I_q} w_i^I(q)\phi_i^I(x)\right] T_q(q'|q, u)\phi(x'; Ax + g(q', u), \mathcal{V}) \\
&\qquad \times \left[\sum_{l=1}^{L} w_{l,n}^{\sigma}(q)\phi(x; \mu_{l,n}^{\sigma}(q), \mathcal{P}_{l,n}^{\sigma}(q))\right] \qquad (62) \\
&= \sum_{q=1}^{N_q}\sum_{l=1}^{L}\sum_{i=1}^{I_q} w_i^I(q)w_{m,n}^{\sigma}(q)Y_q(y^q|q')T_q(q'|q, u)|C^{-1}|\phi(x'; C^{-1}y^x; C^{-1}\mathcal{W}C^{-T}) \\
&\qquad \times \int_{\mathbb{R}^m} \phi(x; \mu_i^I(q), \mathcal{P}_i^I(q))\phi(x; \mu_{l,n}^{\sigma}(q), \mathcal{P}_{l,n}^{\sigma}(q)) \\
&\qquad\qquad \times |A^{-1}|\phi(x; A^{-1}(x' - g(q', u)), A^{-1}\mathcal{V}A^{-T})\, dx \qquad (63) \\
&= \sum_{q=1}^{N_q}\sum_{l=1}^{L}\sum_{i=1}^{I_q} |A^{-1}||C^{-1}|w_i^I(q)w_{m,n}^{\sigma}(q)Y_q(y^q|q')T_q(q'|q, u)\phi(x'; C^{-1}y^x; C^{-1}\mathcal{W}C^{-T}) \\
&\qquad \times \phi(\mu_i^I(q); \mu_{l,n}^{\sigma}(q), \mathcal{P}_i^I(q) + \mathcal{P}_{l,n}^{\sigma}(q)) \int_{\mathbb{R}^m} \phi(x; \tilde{\mu}, \tilde{\mathcal{P}}) \\
&\qquad\qquad \times \phi(x; A^{-1}(x' - g(q', u)), A^{-1}\mathcal{V}A^{-T})\, dx \qquad (64) \\
&= \sum_{q=1}^{N_q}\sum_{l=1}^{L}\sum_{i=1}^{I_q} |A^{-1}||C^{-1}|w_i^I(q)w_{m,n}^{\sigma}(q)Y_q(y^q|q')T_q(q'|q, u)\phi(x'; C^{-1}y^x, C^{-1}\mathcal{W}C^{-T}) \\
&\qquad \times \phi(\mu_i^I(q); \mu_{l,n}^{\sigma}(q), \mathcal{P}_i^I(q) + \mathcal{P}_{l,n}^{\sigma}(q))\phi(A^{-1}(x' - g(q', u)); \tilde{\mu}, \tilde{\mathcal{P}} + A^{-1}\mathcal{V}A^{-T}) \\
&\qquad (65)
\end{aligned}$$

Line (63) follows from (60), line (64) from combining $\phi(x; \mu_i^I(q), \Sigma_i^I(q))$ and $\phi(x; \mu_{l,n}^{\sigma}(q), \mathcal{P}_{l,n}^{\sigma}(q))$ according to (58), and (65) from a second application of (58) and setting the integral of a Gaussian density over $\mathcal{X}$ equal to one.

A final application of (58) and (60) gives

$$\begin{aligned}
\left(\Phi_{y,u}^{g}\sigma_{n,g}\right)(x', q') &= \sum_{q=1}^{N_q}\sum_{l=1}^{L}\sum_{i=1}^{I_q} w_i^I(q)w_{m,n}^{\sigma}(q)Y_q(y^q|q')T_q(q'|q, u)\phi(\mu_i^I(q); \mu_{l,n}^{\sigma}(q), \mathcal{P}_i^I(q) + \mathcal{P}_{l,n}^{\sigma}(q)) \\
&\qquad \times \phi(y^x; C(A\tilde{\mu} + g(q', u)), \mathcal{W} + C\mathcal{V}C^T + CA\tilde{\mathcal{P}}A^TC^T)\phi(x; \mu_{q,l,i,n+1}^{\sigma}(q'), \mathcal{P}_{q,l,i,n+1}^{\sigma}(q'))
\end{aligned}$$

which is again a Gaussian mixture with $N_q L I_q$ components. Specifically,

$$\sigma_{n+1,g}(x',q') = \sum_{q=1}^{N_q}\sum_{l=1}^{L}\sum_{i=1}^{I_q} w^{\sigma}_{q,l,i,n+1}(q')\phi(x';\mu^{\sigma}_{q,l,i,n+1}(q'),\mathcal{P}^{\sigma}_{q,l,i,n+1}(q')),$$

with

$$\begin{aligned}
w^{\sigma}_{q,l,i,n+1}(q') &= w^I_i(q) w^{\sigma}_{l,n}(q) Y_q(y^q|q') T_q(q'|q,u)\phi(\mu^I_i(q);\mu^{\sigma}_{l,n}(q),\mathcal{P}^I_i(q)+\mathcal{P}^{\sigma}_{l,n}(q)) \\
&\quad \times \phi(y^x; C(A\tilde{\mu}+g(q',u)), \mathcal{W}+C\mathcal{V}C^T + CA\tilde{\mathcal{P}}A^TC^T), \\
\mu^{\sigma}_{q,l,i,n+1}(q') &= \mathcal{P}^{\sigma}_{q,l,i,n+1}(q')\left[C^T\mathcal{W}^{-1}y^x + \left(A\tilde{\mathcal{P}}A^T+\mathcal{V}\right)^{-1}(A\tilde{\mu}+g(q,u,q'))\right], \\
\mathcal{P}^{\sigma}_{q,l,i,n+1}(q') &= \left[C^T\mathcal{W}^{-1}C + \left(A\tilde{\mathcal{P}}A^T+\mathcal{V}\right)^{-1}\right]^{-1},
\end{aligned}$$

and

$$\begin{aligned}
\tilde{\mu} &= \tilde{\mathcal{P}}\left[\left(\mathcal{P}^I_i(q)\right)^{-1}\mu^I_i(q) + \left(\mathcal{P}^{\sigma}_{l,n}(q)\right)^{-1}\mu^{\sigma}_{l,n}(q)\right], \\
\tilde{\mathcal{P}} &= \left[\left(\mathcal{P}^I_i(q)\right)^{-1} + \left(\mathcal{P}^{\sigma}_{l,n}(q)\right)^{-1}\right].
\end{aligned}$$

■

Proof of Lemma 7: By construction. Given $\alpha^{*(y)}_{n+1,g}(x',q') = \sum_{m=1}^{M} w^{\alpha,y}_{m,n+1}(q')\phi(x';\mu^{\alpha,y}_{m,n+1}(q'),\mathcal{P}^{\alpha,y}_{m,n+1}(q'))$ for some observation $y \in \mathcal{Y}$, then $\alpha^{g}_{y,u,\sigma}(x,q)$ for control input $u \in \mathcal{U}$ is written

$$\alpha^{g}_{y,u,\sigma}(x,q) = \sum_{q'\in\mathcal{Q}}\int_{\mathbb{R}^m} \alpha^{*(y)}_{n+1,g}(s')\gamma(y|s')\tau(s'|s,u)\,dx' \left[\sum_{i=1}^{I_q} w^I_i(q)\phi^I_i(x)\right] \tag{66}$$

$$\begin{aligned}
&= \sum_{q'\in\mathcal{Q}}\int_{\mathbb{R}^m}\left[\sum_{m=1}^{M} w^{\alpha,y}_{m,n+1}(q')\phi(x';\mu^{\alpha,y}_{m,n+1}(q'),\mathcal{P}^{\alpha,y}_{m,n+1}(q'))\right] \\
&\qquad \times Y_q(y^q|q')\phi(y^x;Cx',\mathcal{W})T_q(q'|q,u) \\
&\qquad\quad \times \phi(x';Ax+g(q',u),\mathcal{V})\,dx'\left[\sum_{i=1}^{I_q} w^I_i(q)\phi^I_i(x)\right]
\end{aligned} \tag{67}$$

$$\begin{aligned}
&= \sum_{q'=1}^{N_q}\sum_{m=1}^{M}\sum_{i=1}^{I_q} w^{\alpha,y}_{m,n+1}(q') w^I_i(q) Y_q(y^q|q') T_q(q'|q,u) \\
&\qquad \times \phi(x;\mu^I_i(q),\mathcal{P}^I_i(q))|C^{-1}|\phi(C^{-1}y^x;\mu^{\alpha,y}_{m,n+1}(q'),C^{-1}\mathcal{W}C^{-T}+\mathcal{P}^{\alpha,y}_{m,n+1}(q')) \\
&\qquad\quad \times \int_{\mathbb{R}^m}\phi(x';\hat{\mu},\hat{\mathcal{P}})\phi(x';Ax+g(q',u),\mathcal{V})\,dx'
\end{aligned} \tag{68}$$

$$\begin{aligned}
&= \sum_{q'=1}^{N_q}\sum_{m=1}^{M}\sum_{i=1}^{I_q} w^{\alpha,y}_{m,n+1}(q') w^I_i(q) Y_q(y^q|q') T_q(q'|q,u) \\
&\qquad \times \phi(x;\mu^I_i(q),\mathcal{P}^I_i(q))\phi(y^x;C\mu^{\alpha,y}_{m,n+1}(q'),\mathcal{W}+C\mathcal{P}^{\alpha,y}_{m,n+1}(q')C^T) \\
&\qquad\quad \times |A^{-1}|\phi(x;A^{-1}(\hat{\mu}-g(q',u)),A^{-1}(\hat{\mathcal{P}}+\mathcal{V})A^{-T})
\end{aligned} \tag{69}$$

$$= \sum_{q'=1}^{N_q} \sum_{m=1}^{M} \sum_{i=1}^{I_q} w_{m,n+1}^{\alpha,y}(q') w_i^I(q) Y_q(y^q|q') T_q(q'|q,u)$$

$$\times \phi(\mu_i^I(q); A^{-1}(\hat{\mu} - g(q',u)), \mathcal{P}_i^I(q) + A^{-1}(\hat{\mathcal{P}} + \mathcal{V})A^{-T})$$

$$\times \phi(y^x; C\mu_{m,n+1}^{\alpha,y}(q'), \mathcal{W} + C\mathcal{P}_{m,n+1}^{\alpha,y}(q')C^T)\phi(x; \mu_{q',m,i,n}^{\alpha,y}(q), \mathcal{P}_{q',m,i,n}^{\alpha,y}(q)) \quad (70)$$

Line (68) follows from one application of (58), line (69) from (60) and another application of (58), and a final product of Gaussian densities gives (70). Hence $\alpha_{y,u,\sigma}^g$ is a Gaussian mixture with $N_q M I_q$ components.

Recalling that $\alpha_{n,g}(x,q) = \sum_{y\in\mathcal{Y}} \alpha_{y,u,\sigma}^g$, it follows that

$$\alpha_{n,g}(x,q) = \sum_{y\in\mathcal{Y}} \sum_{q=1}^{N_q} \sum_{m=1}^{M} \sum_{i=1}^{I_q} w_{y,q',m,i,n+1}^{\sigma}(q) \phi(x; \mu_{y,q',m,i,n+1}^{\alpha}(q), \mathcal{P}_{y,q',m,i,n+1}^{\alpha}(q)),$$

with

$$w_{y,q,l,i,n+1}^{\alpha}(q) = w_{m,n+1}^{\alpha,y}(q') w_i^I(q) Y_q(y^q|q') T_q(q'|q,u)$$

$$\times \phi(\mu_i^I(q); A^{-1}(\hat{\mu} - g(q',u)), \mathcal{P}_i^I(q) + A^{-1}(\hat{\mathcal{P}} + \mathcal{V})A^{-T})$$

$$\times \phi(y^x; C\mu_{m,n+1}^{\alpha,y}(q'), \mathcal{W} + C\mathcal{P}_{m,n+1}^{\alpha,y}(q')C^T)$$

$$\mu_{y,q',m,i,n+1}^{\alpha}(q) = \mathcal{P}_{y,q',m,i,n+1}^{\alpha}(q) \left[\left(\mathcal{P}_i^I(q)\right)^{-1} \mu_i^I(q) + A^T \left(\hat{\mathcal{P}} + \mathcal{V}\right)^{-1} (\hat{\mu} - g(q',u)) \right]$$

$$\mathcal{P}_{y,q',m,i,n+1}^{\alpha}(q) = \left[\left(\mathcal{P}_i^I(q)\right)^{-1} + A^T \left(\hat{\mathcal{P}} + \mathcal{V}\right)^{-1} A \right]^{-1},$$

and

$$\hat{\mu} = \hat{\mathcal{P}} \left[C^T \mathcal{W}^{-1} y^x + \left(\mathcal{P}_{m,n+1}^{\alpha,y}(q')\right)^{-1} \mu_{m,n+1}^{\alpha,y}(q') \right]$$

$$\hat{\mathcal{P}} = \left[C^T \mathcal{W}^{-1} C + \left(\mathcal{P}_{m,n+1}^{\alpha,y}(q')\right)^{-1} \right]^{-1}.$$

■

Proof of Theorem 3: By induction. At time zero, $\sigma_{0,g}(s) = \sigma_0(s)$, so that $\|\sigma_0 - \sigma_{0,g}\|_1 = 0$. Assume that $\|\sigma_i - \sigma_{i,g}\|_1 \leq \gamma_i^\sigma N_q \delta^I$ for all $i = 1, \ldots, n$. Then at time $n+1$ we have, for some $y \in \mathcal{Y}$ and $u \in \mathcal{U}$,

$$\|\sigma_{n+1} - \sigma_{n+1,g}\|_1 \leq \int_{\mathcal{S}} \gamma(y|s') \int_{\mathcal{S}} \left| \mathbf{1}_K(s)\sigma_n(s)\,ds - \sum_{i=1}^{I_q} w_i(q)\phi_i(x)\sigma_{n,g}(s)\,ds \right| \tau(s'|s,u)\,ds'$$

$$\leq \phi_w^* \left[\int_{\mathcal{S}} |\mathbf{1}_K(s)\sigma_n(s)\,ds - \mathbf{1}_K(s)\sigma_{n,g}(s)\,ds| \right.$$

$$\left. + \int_{\mathcal{S}} \left| \mathbf{1}_K(s)\sigma_{n,g}(s)\,ds - \sum_{i=1}^{I_q} w_i(q)\phi_i(x)\sigma_{n,g}(s)\,ds \right| \right]$$

$$\leq \phi_w^* \left[\|\sigma_n - \sigma_{n,g}\|_1 + \sum_{q\in\mathcal{Q}} \int_{\mathbb{R}^m} \left| \mathbf{1}_{K_q}(x) - \sum_{i=1}^{I_q} w_i(q)\phi_i(x) \right| \sigma_{n,g}(x,q)\,dx \right] \quad (71)$$

The first term of line (71) follows because the integral over K is less than the integral over all of $\mathcal{S}$, since K is a compact subset of $\mathcal{S}$. The induction hypothesis completes the proof. ■

Proof of Lemma 8: By induction. At time N,

$$\|\alpha_N - \tilde{\alpha}_{N,g}\|_1 = \sum_{q\in\mathcal{Q}} \int_{\mathbb{R}^m} \left| \mathbf{1}_{K_q}(x) - \sum_{i=1}^{I_q} w_i(q)\phi_i(x) \right| dx \leq N_q \delta^I$$

and the result is satisfied. Assume for $j = N-1, \ldots, n+1$ that $\|\alpha_j^i - \tilde{\alpha}_{j,g}^i\|_1 \leq \left(\sum_{k=j}^N (\lambda\phi_v^*)^{N-k}\overline{\alpha}_{N-k+j+1}\right) N_q\delta^I$ for any $\alpha_j^i \in \Gamma_j$, letting $\overline{\alpha}_{N+1} = 1$. Then for $j = n$,

$$\begin{aligned}
\left\|\alpha_n^i(s) - \tilde{\alpha}_{n,g}^i(s)\right\|_1 = \int_{\mathcal{S}} \Bigg| \int_{\mathcal{S}}\int_{\mathcal{Y}} & \alpha_{n+1}^{i(y)}(s')\gamma(y|s')\tau(s'|s,u^i)\mathbf{1}_K(s)\,dy\,ds' \\
& - \tilde{\alpha}_{n+1,g}^{i(y)}(s')\gamma(y|s')\tau(s'|s,u^i)\sum_{i=1}^{I_q} w_i(q)\phi_i(x) \Bigg| \,dy\,ds'\,ds \\
\leq \int_{\mathcal{S}}\int_{\mathcal{S}}\int_{\mathcal{Y}} & \left|\alpha_{n+1}^{i(y)}(s') - \tilde{\alpha}_{n+1,g}^{i(y)}(s')\right| \gamma(y|s')\tau(s'|s,u^i)\mathbf{1}_K(s)\,dy\,ds'\,ds \\
& + \int_{\mathcal{S}}\int_{\mathcal{S}}\int_{\mathcal{Y}} \tilde{\alpha}_{n+1,g}^{i(y)}(s')\gamma(y|s')\tau(s'|s,u^i) \left|\mathbf{1}_K(s) - \sum_{i=1}^{I_q} w_i(q)\phi_i(x)\right| dy\,ds'\,ds \\
\leq \int_{\mathcal{S}} & \left|\alpha_{n+1}^{i(y)}(s') - \tilde{\alpha}_{n+1,g}^{i(y)}(s')\right| \phi_v^* ds' \int_{\mathcal{S}} \mathbf{1}_K(s) ds + \sum_{q\in\mathcal{Q}}\int_{\mathbb{R}^m} \overline{\alpha}_{n+1} \left|\mathbf{1}_{K_q}(x) - \sum_{i=1}^{I_q} w_i(q)\phi_i(x)\right| dx \\
\leq & \left(\sum_{k=n}^N (\lambda\phi_v^*)^{N-k}\overline{\alpha}_{N-k+n+1}\right) N_q\delta^I
\end{aligned}$$

■

Proof of Theorem 4: By construction. For any time $n \in [0, N]$, given $\sigma_n \in \Sigma_n$ and $\sigma_{n,g} \in \Sigma_{n,g}$, we can rewrite the value function evaluated at σ in terms of α-functions.

$$\begin{aligned}
\left|V_n^*(\sigma_n) - V_{n,g}^*(\sigma_{n,g})\right| &= \left| \sup_{\alpha_n\in\Gamma_n} \langle \alpha_n, \sigma_n\rangle - \sup_{\alpha_{n,g}\in\Gamma_{n,g}} \langle \alpha_{n,g}, \sigma_{n,g}\rangle \right| \\
&= \left|\langle \alpha_n^k, \sigma_n\rangle - \langle \alpha_{n,g}^l, \sigma_{n,g}\rangle\right|
\end{aligned}$$

As in the proof of Theorem 2, assume without loss of generality that $\langle \alpha_n^k, \sigma_n\rangle \geq \langle \alpha_{n,g}^l, \sigma_{n,g}\rangle$.

$$\begin{aligned}
\left|V_n^*(\sigma_n) - V_{n,g}^*(\sigma_{n,g})\right| &= \langle \alpha_n^k, \sigma_n\rangle - \langle \alpha_{n,g}^l, \sigma_{n,g}\rangle \\
&\leq \langle \alpha_n^k, \sigma_n\rangle - \langle \tilde{\alpha}_{n,g}^k, \sigma_{n,g}\rangle \\
&\leq \left|\langle \alpha_n^k, \sigma_n\rangle - \langle \tilde{\alpha}_{n,g}^k, \sigma_n\rangle\right| + \left|\langle \tilde{\alpha}_{n,g}^k, \sigma_n\rangle - \langle \tilde{\alpha}_{n,g}^k, \sigma_{n,g}\rangle\right| \\
&\leq \int_{\mathcal{S}} \left|\alpha_n^k(s) - \tilde{\alpha}_{n,g}^k(s)\right| \sigma_n(s)\,ds + \int_{\mathcal{S}} \tilde{\alpha}_{n,g}^k(s)\left|\sigma_n(s) - \sigma_{n,g}(s)\right| ds \\
&\leq \left(\sum_{k=n+1}^N (\lambda\phi_v^*)^{N-k}\overline{\alpha}_{N-k+n+1}\right)\phi_{\sigma,n}^* N_q\delta^I + \overline{\alpha}_n\gamma_n^\sigma N_q\delta^I
\end{aligned}$$

■

Proof of Lemma 10: Define $\overline{K}$ such that $Y_x(\mathcal{Y}\backslash\overline{K}|z \in K_\delta, u) < \frac{\epsilon}{N}$. Then

$$\begin{aligned}
V^*_{n,\delta}(\sigma_{n,\delta}) - \hat{V}^*_{n,\delta}(\sigma_{n,\delta}) &= \sup_{\alpha_{n,\delta}\in\Gamma_{n,\delta}} \sum_{z\in K_\delta} \alpha_{n,\delta}(z)\sigma_{n,\delta}(z) - \sup_{\hat{\alpha}\in\tilde{\Gamma}_{n,\delta}} \sum_{z} \hat{\alpha}_{n,\delta}(z)\sigma_{n,\delta}(z) \\
&\leq \int_{\overline{K}} \sum_{z,z'\in K_\delta} \Big[\alpha^{*(y)}_{n+1,\delta}(z')\gamma(y|z')\tau_\delta(z'|z,u^*)\sigma_{n,\delta}(z)\,dy \\
&\qquad -\alpha^{*(\theta(y))}_{n+1,\delta}(z')\gamma(y|z')\tau_\delta(z'|z,u^*)\sigma_{n,\delta}(z)\,dy\Big] \\
&\qquad + \int_{\mathcal{Y}\backslash\overline{K}} \sum_{z,z'\in K_\delta} \alpha^{*(y)}_{n+1,\delta}(z')\gamma(y|z')\tau_\delta(z'|z,u^*)\sigma_{n,\delta}(z)\,dy \\
&\qquad - \int_{\mathcal{Y}\backslash\overline{K}} \sum_{z,z'\in K_\delta} \alpha^{*(\psi_y)}_{n+1,\delta}(z')\gamma(y|z')\tau_\delta(z'|z,u^*)\sigma_{n,\delta}(z)\,dy \\
&\leq \int_{\overline{K}} \sum_{z,z'\in K_\delta} \Big[\alpha^{*(y)}_{n+1,\delta}(z')\gamma(y|z')\tau_\delta(z'|z,u^*)\sigma_{n,\delta}(z)\,dy \\
&\qquad -\alpha^{*(\theta(y))}_{n+1,\delta}(z')\gamma(y|z')\tau_\delta(z'|z,u^*)\sigma_{n,\delta}(z)\,dy\Big] + \frac{\epsilon}{N} \qquad (72)
\end{aligned}$$

Note that (72) is nonnegative, meaning that using $\hat{\alpha}^*_\delta$ produces a lower bound to the actual value function given by α^*_δ. This follows because $\alpha^{*(y)}_{n+1,\delta}$ is chosen optimally for only a subset of $\mathcal{Y}$ (at the points $\theta(y)$), and for all other $y \in \mathcal{Y}$, $\alpha_{n+1,\delta}$ is suboptimal, producing a lower value.

Next, we can bound $\alpha^{*(\theta(y))}_{n+1,\delta}\gamma(y|z')$ from below based on how the points w are defined.

$$\begin{aligned}
V^*_{n,\delta}(\sigma_\delta) - \hat{V}^*_{n,\delta}(\sigma_{n,\delta}) &\leq \int_{\overline{K}} \left[\alpha^{*(y)}_{n+1,\delta}(z')\gamma(y|z') - \alpha^{*(\theta(y))}_{n+1,\delta}(z')\gamma(\theta(y)|z')\right]\,dy + \frac{\epsilon}{N} \\
&\leq N_q\overline{\lambda}h^{(1)}_y\delta^y + \frac{\epsilon}{N}
\end{aligned}$$

■

Proof of Theorem 6: By induction. At time N, $V^*_{N,\delta}(\sigma_{n,\delta}) = \tilde{V}^*_{N,\delta}(\sigma_{n,\delta})$ since $\Gamma_N = \tilde{\Gamma}_N = \mathbf{1}_{K_\delta}$. Assume for all $i = n+1, \ldots, N-1$ that $V^*_{i,\delta}(\sigma_{n,\delta}) - \tilde{V}^*_{i,\delta}(\sigma_{n,\delta}) \leq (N-i)N_q\overline{\lambda}h^{(1)}_y\delta^y + \frac{(N-j)\epsilon}{N}$. Then, at time n,

$$\begin{aligned}
V^*_{n,\delta}(\sigma_{n,\delta}) - \tilde{V}^*_{n,\delta}(\sigma_{n,\delta}) &= \langle\alpha^*_{n,\delta}, \sigma_{n,\delta}\rangle - \langle\tilde{\alpha}^*_{n,\delta}, \sigma_{n,\delta}\rangle \\
&= \langle\alpha^*_{n,\delta}, \sigma_{n,\delta}\rangle - \langle\hat{\alpha}^*_{n,\delta}, \sigma_{n,\delta}\rangle + \langle\hat{\alpha}^*_{n,\delta}, \sigma_{n,\delta}\rangle - \langle\tilde{\alpha}^*_{n,\delta}, \sigma_{n,\delta}\rangle \\
&\leq N_q\overline{\lambda}h^{(1)}_y\delta^y + \int_{\mathcal{Y}} \sup_{\Gamma_{n+1,\delta}} \langle\alpha_{n+1,\delta}, \Phi^\delta_{y,\hat{u}^*}\sigma_{n,\delta}\rangle\mathbb{P}_\delta(dy|\sigma_{n,\delta}, \hat{u}^*) \\
&\qquad - \int_{\mathcal{Y}} \sup_{\tilde{\Gamma}_{n+1,\delta}} \langle\tilde{\alpha}_{n+1,\delta}, \Phi^\delta_{y,\tilde{u}^*}\sigma_\delta\rangle\mathbb{P}_\delta(dy|\sigma_{n,\delta}, \tilde{u}^*) + \frac{\epsilon}{N} \\
&\leq N_q\overline{\lambda}h^{(1)}_y\delta^y + V^*_{n+1,\delta}(\sigma_{n+1,\delta}) - \tilde{V}^*_{n+1,\delta}(\sigma_{n+1,\delta}) + \frac{\epsilon}{N}
\end{aligned}$$

Applying the induction inequality and combining terms completes the proof. ■

REFERENCES

[1] C. Tomlin, I. Mitchell, A. Bayen, and M. Oishi, "Computational techniques for the verification and control of hybrid systems," in *Proceedings of the IEEE*, vol. 91, no. 7, July 2003, pp. 986–1001.

[2] M. Prandini and J. Hu, *Stochastic Reachability: Theoretical Foundations and Numerical Approximation*, ser. Lecture Notes in Control and Information Sciences. Springer Verlag, 2006, pp. 107–139.

[3] I. Mitchell and J. Templeton, "A toolbox of Hamilton-Jacobi solvers for analysis of nondeterministic continuous and hybrid systems," in *Hybrid Systems: Computation and Control*, 2005, vol. 3414, pp. 480–494.

[4] A. Abate, M. Prandini, J. Lygeros, and S. Sastry, "Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems," *Automatica*, vol. 44, no. 11, pp. 2724–2734, 2008.

[5] S. Summers and J. Lygeros, "Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem," *Automatica*, vol. 46, no. 12, pp. 1951–1961, 2010.

[6] R. Verma and D. del Vecchio, "Safety control of hidden mode hybrid systems," *IEEE Transactions on Automatic Control*, vol. 57, no. 1, pp. 62–77, 2012.

[7] R. Ghaemi and D. del Vecchio, "Control for safety specifications of systems with imperfect information on a partial order," *IEEE Transactions on Automatic Control*, vol. 59, no. 4, pp. 982 – 995, 2014.

[8] J. Ding, A. Abate, and C. Tomlin, "Optimal control of partially observable discrete time stochastic hybrid systems for safety specifications," in *American Control Conference*, 2013, pp. 6231–6236.

[9] K. Lesser and M. Oishi, "Reachability for partially observable discrete time stochastic hybrid systems," *Automatica*, vol. 50, no. 8, pp. 1989–1998, 2014.

[10] ——, "Finite state approximation for verification of partially observable stochastic hybrid systems," in *Hybrid Systems: Computation and Control*, 2015.

[11] D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 2005, vol. 1.

[12] A. Abate, J.-P. Katoen, J. Lygeros, and M. Prandini, "Approximate model checking of stochastic hybrid systems," *European Journal of Control*, vol. 16, no. 6, pp. 624–641, 2010.

[13] S. Soudjani and A. Abate, "Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes," *SIAM Journal on Applied Dynamical Systems*, vol. 12, no. 2, pp. 921–956, 2013.

[14] A. A. Julius and G. J. Pappas, "Approximations of stochastic hybrid systems," *IEEE Transactions on Automatic Control*, vol. 54, no. 6, pp. 1193 – 1203, 2009.

[15] M. Fränzle, E. M. Hahn, H. Hermanns, N. Wolovick, and L. Zhang, "Measurability and safety verification for stochastic hybrid systems," in *Hybrid Systems: Computation and Control*, 2011, pp. 43–52.

[16] M. Zamani, I. Tkachev, and A. Abate, "Bisimilar symbolic models for stochastic control systems without state-space discretization," in *Hybrid Systems: Computation and Control*, 2014, pp. 41 – 50.

[17] C. Lusena, J. Goldsmith, and M. Mundhenk, "Nonapproximability results for partially observable Markov decision processes," *Journal of Artificial Intelligence Research*, vol. 14, pp. 83–103, 2001.

[18] G. Shani, J. Pineau, and R. Kaplow, "A survey of point-based POMDP solvers," *Autonomous Agents and Multi-Agent Systems*, vol. 27, no. 1, pp. 1–51, 2013.

[19] E. Sondik, "The optimal control of partially observable Markov processes," Ph.D. dissertation, Stanford University, 1971.

[20] A. Brooks, A. Makarenko, S. Williams, and H. Durrant-Whyte, "Parametric POMDPs for planning in continuous state spaces," *Robotics and Autonomous Systems*, vol. 54, no. 11, pp. 887–897, 2006.

[21] E. Zhou, M. Fu, and S. Marcus, "Solving continuous-state POMDPs via density projection," *IEEE Transactions on Automatic Control*, vol. 55, no. 5, pp. 1101–1116, 2010.

[22] J. M. Porta, N. Vlassis, M. T. Spain, and P. Poupart, "Point-based value iteration for continuous POMDPs," *Journal of Machine Learning Research*, vol. 7, pp. 2329–2367, 2006.

[23] E. Brunskill, L. Kaelbling, T. Lozano-Perez, and N. Roy, "Planning in partially-observable switching-mode continuous domains," *Annals of Mathematics and Artificial Intelligence*, vol. 58, pp. 185–216, 2010.

[24] D. P. Bertsekas and S. E. Shreve, *Stochastic Optimal Control: The Discrete-Time Case*. Athena Scientific, 1996.

[25] J. Pineau, G. Gordon, and S. Thrun, "Anytime point-based approximations for large POMDPs," *Journal of Artificial Intelligence Research*, vol. 27, pp. 335–380, 2006.

[26] K. Lesser and M. Oishi, "Approximate verification of partially observable discrete time stochastic hybrid systems," 2014, preprint, http://arxiv.org/abs/1410.8054.

[27] J. Park and I. Sandberg, "Universal approximation using radial-basis-function networks," *Neural Computation*, vol. 3, no. 2, pp. 246 – 257, 1991.

[28] B. Fornberg and N. Flyer, "The Gibbs phenomenon for radial basis functions," in *Gibbs Phenomenon in Various Representations and Applications*, 2011, pp. 201–224.

[29] K. Zhang and J. Kwok, "Simplifying mixture models through function approximation," *IEEE Transactions on Neural Networks*, vol. 21, no. 4, pp. 644–658, 2010.

[30] H. Kurniawati, D. Hsu, and W. S. Lee, "SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces," in *Robotics: Science and Systems*, 2008.